3-connected Reduction for Regular Graph Covers☆,☆☆

Jiří Fiala[b], Pavel Klavík[a], Jan Kratochvíl[b], Roman Nedela[c,d]

[a]*Computer Science Institute, Faculty of Mathematics and Physics, Charles University, Malostranské náměstí 25, 118 00 Prague, Czech Republic.*
[b]*Department of Applied Mathematics, Faculty of Mathematics and Physics, Charles University in Prague, Malostranské náměstí 25, 118 00 Prague, Czech Republic.*
[c]*Institute of Mathematics and Computer Science SAS, Ďumbierska 1, 974 11 Banská Bystrica, Slovak republic.*
[d]*European Centre of Excellence NTIS, University of West Bohemia, Pilsen, Czech Republic.*

Abstract

A graph G *covers* a graph H if there exists a locally bijective homomorphism from G to H. We deal with *regular covers* in which this homomorphism is prescribed by an action of a semiregular subgroup Γ of $\mathrm{Aut}(G)$; so $H \cong G/\Gamma$. In this paper, we study behaviour of regular graph coverings with respect to 1-cuts and 2-cuts in G.

We describe *reductions* which produce a series of graphs $G = G_0, \dots, G_r$ such that G_{i+1} is created from G_i by replacing certain inclusion minimal subgraphs with colored edges. The process ends with a *primitive graph* G_r which is either 3-connected, or a cycle, or K_2. This reduction can be viewed as a non-trivial modification of reductions of Trachtenbrot (1958), Tutte (1966), Hopcroft and Tarjan (1973), Cuningham and Edmonds (1980), and Walsh (1982). A novel feature of our approach is that in each step all essential information about symmetries of G are preserved.

A regular covering $G_0 \to H_0$ induces regular coverings $G_i \to H_i$ where H_i is the i-th reduction of H_0. This property allows to construct all possible quotients H_0 of G_0 from the possible quotients H_r of G_r. By applying this method to planar graphs, we give a direct proof of Negami's Theorem (1988). Our structural results are also used in subsequent papers for regular covering testing when G is a planar graph and for characterization of the automorphism groups of planar graphs.

Keywords: regular graph covers, 3-connected reduction, quotient expansion, half-quotients

1. Introduction

The notion of *covering* originates in topology to describe local similarity of two topological spaces. In this paper, we study coverings of graphs in a more restricting version called *regular covering*, for which the covering projection is described by a semiregular action of a group; see Section 2 for the formal definition. If G regularly covers H, then H is called a *quotient* of G. See Fig. 1 for an example.

Regular graph covers have many applications in graph theory, for instance they were used to solve the Heawood map coloring problem [25, 16] and to construct arbitrarily large highly symmetrical graphs [5]. The concept of a regular covering of graphs gives rise to a powerful construction of large graphs with prescribed properties from smaller ones. It can be demonstrated by the well-known construction of a Cayley graph, where a large graph is defined by specifying few generators of a group. Regular covers are generalizations of Cayley graphs which are the graphs having one-vertex regular quotient.

☆This paper continues the research started in ICALP 2014 [13] and extends its results. For a structural diagram visualizing our results, see `http://pavel.klavik.cz/orgpad/regular_covers.html` (supported for Firefox and Google Chrome).

☆☆This work was initiated during workshops Algebraic, Topological and Complexity Aspects of Graph Covers (ATCACG). The authors are supported by CE-ITI (P202/12/G061 of GAČR). The first author is also supported by the project Kontakt LH12095, the second and the third authors by Charles University as GAUK 196213, the fourth author by the Ministry of Education of the Slovak Republic, the grant VEGA 1/0150/14, by Project LO1506 of the Czech Ministry of Education and by the project APVV-15-0220.

Email addresses: `fiala@kam.mff.cuni.cz` (Jiří Fiala), `klavik@iuuk.mff.cuni.cz` (Pavel Klavík), `honza@kam.mff.cuni.cz` (Jan Kratochvíl), `nedela@savbb.sk` (Roman Nedela)

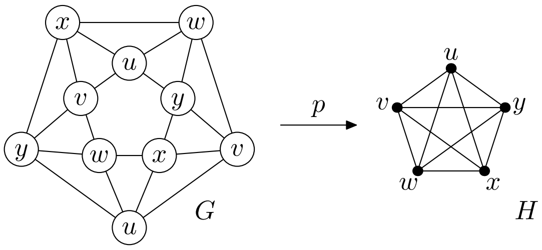

Figure 1: A regular covering projection p from a graph G to one of its quotients H. For every vertex $v \in \boldsymbol{V}(G)$, the image $p(v)$ is written in the circle.

1.1. Our Results

We fully describe the behaviour of regular covering with respect to 1-cuts and 2-cuts in G. It is closely related to the behaviour of 1- and 2-cuts under a semiregular action of a subgroup of the automorphism group $\mathrm{Aut}(G)$ of G. For 1-cuts, it is quite simple since $\mathrm{Aut}(G)$ fixes the central block. But the behaviour of regular covering on 2-cuts is complex. Our main result Theorem 1.2 describes all possible quotients of some graph class if we understand quotients of 3-connected graphs in this class. Our result can be applied to planar graphs and its consequences are described in Sections 1.2 and 1.3.

We process the graph G by a series of *reductions* replacing some parts of G, called *atoms*, separated by 1- and 2-cuts, by edges. This very natural idea of reductions was first introduced in a seminal paper of Trakhtenbrot [28] and further extended in [29, 18, 10, 30]. Our reduction has two key differences:

- The papers [28, 29, 18, 10, 30] apply the reduction only to 2-connected graphs. We also reduce parts separated by 1-cuts, so we generalize the definition of atoms. We need to do this since a quotient of a 2-connected graph might only be 1-connected, so the unified treatment is desirable.
- The reduction is augemented by colored edges (encoding different isomorphism classes) of three different types (encoding different symmetry types of atoms). This allows to capture the changes in the automorphism group and its semiregular subgroups.

Atoms and Reductions. In Section 3, we introduce the important definition of an *atom*. Atoms are inclusion-minimal subgraphs with respect to 1-cuts and 2-cuts which cannot be further simplified and are essentially paths, cycles, stars or 3-connected. The reduction constructs a series of graphs $G = G_0, G_1, \dots, G_r$. The reduction from G_i to G_{i+1} is done by replacing all the atoms of G_i by colored edges, where the colors encode the isomorphism classes of atoms. In this process, we remove details from the graph but preserve its overall structure. Our definition of atoms is quite technical, so that it works nicely with respect to regular covering as it will become clear later. The last (irreducible) graph in the sequence, denoted by G_r, is called *primitive*. It is either very simple (K_2 or a cycle), or 3-connected. Therefore, we call this reduction process the 3-*connected reduction* of G.

When the graph G is not 3-connected, we consider its block-tree. The central block plays the key role in every regular covering projection. The reason is that a covering $G \to H$ behaves non-trivially only on this central block; the remaining blocks are isomorphically preserved in H. Therefore the atoms are defined with respect to the central block. We distinguish three types of atoms:

- *Proper atoms* are inclusion-minimal subgraphs separated by a 2-cut inside a block.
- *Dipoles* are formed by the sets of all parallel edges joining two vertices.
- *Block atoms* are blocks which are leaves of the block-tree, or stars of all pendant edges attached to a vertex. The central block is never a block atom.

The key properties of the automorphism groups are preserved by the reductions. More precisely, the reduction from G_i to G_{i+1} is defined in a way that an induced *reduction epimorphism* $\Phi_i : \mathrm{Aut}(G_i) \to \mathrm{Aut}(G_{i+1})$ possesses nice properties; see Proposition 4.2. Using it, we can describe the change of the automorphism group explicitly:

Proposition 1.1. *If G_i is reduced to G_{i+1}, then*

$$\mathrm{Aut}(G_{i+1}) \cong \mathrm{Aut}(G_i)/\mathrm{Ker}(\Phi_i).$$

Expansions. We aim to investigate how the knowledge of regular quotients of G_{i+1} can be used to construct all regular quotients of G_i. To do so, we introduce the reversal of the reduction called the *expansion*. If $H_{i+1} = G_{i+1}/\Gamma_{i+1}$, then the expansion produces H_i by replacing colored edges back by atoms. To do this, we have to understand how regular covering behaves with respect to atoms. Inspired by Negami [24], we show that each proper atom/dipole has three possible types of quotients that we call an *edge-quotient*, a *loop-quotient* and a *half-quotient*. The edge-quotient and the loop-quotient are uniquely determined but an atom may have many non-isomorphic half-quotients.

The constructed quotients contain colored edges, loops and half-edges corresponding to atoms. Each half-edge in H_{i+1} is created from a halvable edge if an automorphism of Γ_{i+1} fixes this halvable edge and exchanges its endpoints. Roughly speaking it corresponds to cutting the edge in half. The following theorem is our main result and it describes every possible expansion of H_{i+1} to H_i:

Theorem 1.2. *Let G_{i+1} be a reduction of G_i. Every quotient H_i of G_i can be constructed from some quotient H_{i+1} of G_{i+1} by replacing each edge, loop and half-edge of H_{i+1} by the subgraph corresponding to the edge-, the loop-, or a half-quotient of an atom of G_i, respectively.*

Suppose that some regular quotient of the primitive graph G_r is chosen, so $H_r = G_r/\Gamma_r$. The above theorem allows to describe all regular quotients H of G rising from H_r, as depicted in the diagram in Fig. 2.

1.2. Algorithmic and Complexity Consequences

Our main algorithmic motivation is the study of the computational complexity of regular covering testing:

Problem:	RegularCover
Input:	Connected graphs G and H.
Question:	Does G regularly cover H?

Our structural results have the following algorithmic implications, described in [13, 14]:

Theorem 1.3 (Fiala et al. [13, 14]). *If G is planar, we can solve* RegularCover *in time $\mathcal{O}(n^c \cdot 2^{\mathbf{e}(H)/2})$, where c is a constant and $\mathbf{e}(H)$ is the number of edges of H.*

Theorem 1.2 suggests that there might be exponentially many quotients of G, and so this algorithm has to test efficiently whether H is one of them. In particular, for every fixed graph H, the constructed algorithm runs in polynomial time.

Relations to General Covers. The aforementioned decision problem is closely related to the complexity of general covering testing which was widely studied before. We try to understand how much the additional algebraic structure influences the computational complexity. Study of the complexity of general covers was pioneered by Bodlaender [7] in the context of networks of processors in parallel computing, and for fixed target graph was first asked by Abello et al. [1]. The problem H-Cover asks whether an input graph G covers a fixed graph H. The general complexity is still unresolved but papers [21, 12] show that it is NP-complete for every r-regular graph H where $r \geq 3$. For a survey, see [15].

The complexity results concerning graph covers are mostly NP-complete. In our impression, the additional algebraic structure of regular graph covers makes the problem easier, as shown by the following two contrasting results. The problem H-Cover remains NP-complete for several small fixed graphs H (such as K_4, K_5) even for planar inputs G [6]. On the other hand, Theorem 1.3 shows that for planar graphs G the problem RegularCover is fixed-parameter tractable in the number of edges of H.

reduction

$$\begin{tikzcd}
G = G_0 \arrow[r] \arrow[d, "\Gamma_0"] & G_1 \arrow[r] \arrow[d, "\Gamma_1"] & \cdots \arrow[r] & G_i \arrow[r] \arrow[d, "\Gamma_i"] & G_{i+1} \arrow[r] \arrow[d, "\Gamma_{i+1}"] & \cdots \arrow[r] & G_r \arrow[d, "\Gamma_r"] \\
H_0 & H_1 \arrow[l] & \cdots \arrow[l] & H_i \arrow[l] & H_{i+1} \arrow[l] & \cdots \arrow[l] & H_r \arrow[l]
\end{tikzcd}$$

expansion

Figure 2: The reduction is on top, the expansion is on bottom. It holds that $H_i = G_i/\Gamma_i$ and Γ_i is a group extension of Γ_{i+1}.

Relations to Cayley Graphs and Graph Isomorphism. The notion of regular graph covers builds a bridge between two seemingly different problems. If the base graph H is a one-vertex graph, it corresponds to the problem of recognition of Cayley graphs whose complexity is not known. A polynomial-time algorithm is known only for circulant graphs [11]. When both graphs G and H have the same size, we get graph isomorphism testing. Our results are far from solving these problems, but we believe that better understanding of regular covering can also shed new light on these famous problems.

Theoretical motivation for studying graph isomorphism is very similar to RegularCover. For practical instances, one can solve the isomorphism problem very efficiently using various heuristics. But a polynomial-time algorithm working for all graphs is not known and it is very desirable to understand the complexity of graph isomorphism. It is known that testing graph isomorphism is equivalent to testing isomorphism of general mathematical structures [17]. The notion of isomorphism is widely used in mathematics when one wants to show that two seemingly different mathematical structures are the same. One proceeds by guessing a mapping and proving that this mapping is an isomorphism. The natural complexity question is whether there is a better way in which one algorithmically derives an isomorphism. Similarly, regular covering is a well-known mathematical concept which is algorithmically interesting and not understood.

Further, a regular covering is described by a semiregular subgroup of the automorphism group $\mathrm{Aut}(G)$. Therefore it seems to be closely related to computation of $\mathrm{Aut}(G)$, since one should have a good understanding of this group first, to solve the regular covering problem. The problem of computing automorphism groups is known to be closely related to graph isomorphism.

1.3. Structural Consequences

Direct Proof of Negami Theorem. In 1988, Seiya Negami [24] proved that a connected graph H has a finite regular planar cover G if and only if H is projective planar. If the graph G is 3-connected, then $\mathrm{Aut}(G)$ is a *spherical group*. Therefore the conjecture can be easily proved using geometry, since a quotient of the sphere is either the disk, the sphere, or the projective plane. The hard part of the proof is to deal with graphs G containing 1-cuts and 2-cuts.

Negami considered a minimal counterexample. In his proof, an essence of the crucial notion of an atom appears. A regular covering projection can behave on an atom in three different ways, and this understanding can be used to make the minimal counterexample smaller which forces a contradiction. In comparison, our work goes further and structurally describes all possible quotients H of a planar graph G. It is well known that quotients of 3-connected planar graphs can be described geometrically; see Section 5. Therefore, Theorem 1.2 describes all quotients of planar graphs.

Characterizing Automorphism Groups of Planar Graphs. We can use the key property that our reductions preserve essential information about symmetries of the graph G to describe the automorphism groups of planar graphs. This shows that computing automorphism groups can be reduced to computing them for 3-connected graphs which are the spherical groups. The automorphism groups of planar graphs were determined by Babai [2, 3] using a similar approach. An idea how to compute a generating set of these groups was described by Colbourn and Booth [8], but it was never fully developed.

Our reduction is used in [20] to present a clear structural description of automorphism groups of planar graphs. First, Proposition 1.1 is strengthened for planar graphs to show that $\mathrm{Aut}(G_i) = \mathrm{Aut}(G_{i+1}) \ltimes \mathrm{Ker}(\Phi_i)$. After that the following inductive characterization of stabilizers of vertices in connected planar graphs $\mathrm{Fix}(\mathsf{PLANAR})$ is described. It is similar to Jordan's characterization of the automorphism groups of trees [19]:

Theorem 1.4 (Klavík et al. [20]). *The class of groups* $\mathrm{Fix}(\mathsf{PLANAR})$ *is defined inductively as follows:*

(a) $\{1\} \in \mathrm{Fix}(\mathsf{PLANAR})$.
(b) If $\Psi_1, \Psi_2 \in \mathrm{Fix}(\mathsf{PLANAR})$, *then* $\Psi_1 \times \Psi_2 \in \mathrm{Fix}(\mathsf{PLANAR})$.
(c) If $\Psi \in \mathrm{Fix}(\mathsf{PLANAR})$, *then* $\Psi \wr \mathbb{S}_n, \Psi \wr \mathbb{C}_n \in \mathrm{Fix}(\mathsf{PLANAR})$.
(d) If $\Psi_1, \Psi_2, \Psi_3 \in \mathrm{Fix}(\mathsf{PLANAR})$, *then* $(\Psi_1^{2n} \times \Psi_2^n \times \Psi_3^n) \rtimes \mathbb{D}_n \in \mathrm{Fix}(\mathsf{PLANAR})$.
(e) If $\Psi_1, \Psi_2, \Psi_3, \Psi_4, \Psi_5, \Psi_6 \in \mathrm{Fix}(\mathsf{PLANAR})$, *then*

$$(\Psi_1^4 \times \Psi_2^2 \times \Psi_3^2 \times \Psi_4^2 \times \Psi_5^2 \times \Psi_6) \rtimes \mathbb{C}_2^2 \in \mathrm{Fix}(\mathsf{PLANAR}).$$

Next, the class of automorphism groups of connected planar graphs, denoted $\mathrm{Aut}(\text{connected } \mathsf{PLANAR})$, are characterized as follows:

Theorem 1.5 (Klavík et al. [20]). *The class* Aut(connected PLANAR) *consists of the following groups. Let G' be a planar graph with colored vertices and colored (possibly oriented) edges, which is either 3-connected, or K_2, or a cycle C_n. Let $m_1, \dots, m_\ell$ be the sizes of the vertex- and edge-orbits of the action of* $\mathrm{Aut}(G')$. *Then for any choices $\Psi_1, \dots, \Psi_\ell \in \mathrm{Fix}(\mathsf{PLANAR})$, we have*

$$(\Psi_1^{m_1} \times \cdots \times \Psi_\ell^{m_\ell}) \rtimes \mathrm{Aut}(G') \in \mathrm{Aut}(\text{connected } \mathsf{PLANAR}).$$

Every group of Aut(connected PLANAR) *can be constructed in the above way.*

This characterization leads to a quadratic-time algorithm for computing these automorphism groups.

2. Definitions and Preliminaries

2.1. Model of Graph

The concept of graph covering comes from topological graph theory where graphs are understood as 1-dimensional CW-complexes. This means that edges are represented by real open intervals, vertices are points, and the topological closure of an edge e is either a closed interval, or a simple cycle. In the first case, e joins two different vertices u and v incident to e. In the second case, e is incident just to one vertex v and e is a loop based at v. When one considers regular quotients of graphs, a third type of edges may appear [22]. For a non-trivial involution swapping the end-vertices of an edge e, the regular covering projection maps e to an "edge" whose one end is incident to a vertex while the other is free. Its topological closure is homeomorphic to a half-closed interval, and we call it a half-edge.

We use this extended definition of graphs, and we describe it for a non-familiar reader. A *graph* G is a tuple $(\boldsymbol{H}, \boldsymbol{V}, \iota, \lambda)$, where $\boldsymbol{H}$ is a set of half-edges, $\boldsymbol{V}$ is a set of vertices, $\iota : \boldsymbol{H} \to \boldsymbol{V}$ is a partial function of incidence, and $\lambda : \boldsymbol{H} \to \boldsymbol{H}$ is an involution, pairing half-edges. The set of edges $\boldsymbol{E}$ is formed by orbits of λ of size 2, while orbits of size 1 form standalone half-edges. Each edge $\{h, \lambda h\}$ is one of the four kinds: a *standard edge* if $\iota(h) \neq \iota(\lambda h)$, a *loop* if $\iota(h) = \iota(\lambda h)$, a *pendant edge* where exactly one of $\iota(h)$ and $\iota(\lambda h)$ is not defined, and a *free edge* where both $\iota(h)$ and $\iota(\lambda h)$. For (standalone) half-edges h, we have $h = \lambda h$ and it is called a *free half-edge* when $\iota(h) = \iota(\lambda h)$ is not defined. Standalone half-edges are mostly called half-edges and we do not distinguish between h and the orbit $\{h\}$ of size 1.[1] See Fig. 3a.

Unless the graph is K_2, we remove all vertices of degree 1 while keeping both half-edges (one with ι not defined). Assuming that the original graph contains no pendant edges, this removal does not change the automorphism group and existence of regular covering projections from G to H (when the removal is applied on both G and H). A pendant edge attached to v is called a *single pendant edge* if it is the only pendant edge attached to v. Most graphs in this paper are assumed to be connected, so they contain no free edges and half-edges. (Sometimes we consider subgraphs which may be disconnected and may contain them.)

When we work with several graphs, we use $\boldsymbol{H}(G)$, $\boldsymbol{V}(G)$, and $\boldsymbol{E}(G)$ to denote the sets of half-edges, vertices and edges of G, respectively. We denote $|\boldsymbol{H}(G)|$ by $\boldsymbol{h}(G)$, $|\boldsymbol{V}(G)|$ by $\boldsymbol{v}(G)$, $|\boldsymbol{E}(G)|$ by $\boldsymbol{e}(G)$. When

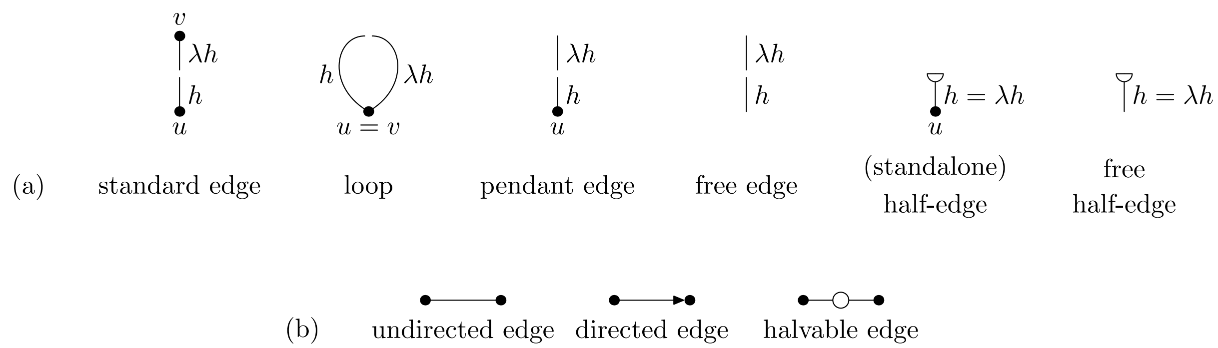

Figure 3: (a) Four kinds of edges and two kinds of half-edges are depicted. We highlight two half-edges composing each edge by a small gap, omitted in the remaining figures. To distinguish pendant edges from standalone half-edges, we end the latter by a half-circle. (b) Three possible types for standard edges and loops (pendant edges are always undirected). We note that only halvable edges may be projected to standalone half-edges which corresponds to cutting the middle circle in half, explaining the symbol for half-edges.

[1] When the distinction is needed, some papers call the elements of $\boldsymbol{H}$ as darts or arcs while standalone half-edges are called half-edges or semiedges.

G contains no standalone half-edges, clearly $\boldsymbol{h}(G) = 2\boldsymbol{e}(G)$. Also, we consider graphs with colored edges and also with three different edge types (directed edges, undirected edges and a special type called halvable edges, see Fig. 3b). It might seem strange to consider such general objects. But when we apply reductions, we replace parts of the graph by edges and the colors encode isomorphism classes of replaced parts. Even if G and H are simple, the more general colored multigraphs are naturally constructed in the process of reductions.

2.2. Automorphism Groups

Automorphisms. We state the definitions in a very general setting of multigraphs with half-edges. An *automorphism* π is fully described by a permutation $\pi_h : \boldsymbol{H}(G) \to \boldsymbol{H}(G)$ preserving edges and incidences between half-edges and vertices, i.e., $\pi_h(\lambda h) = \lambda \pi_h(h)$ and $\pi_h(\iota(h)) = \iota(\pi_h(h))$ (where either both, or none are defined). Thus, π_h induces two permutations $\pi_v : \boldsymbol{V}(G) \to \boldsymbol{V}(G)$ and $\pi_e : \boldsymbol{E}(G) \to \boldsymbol{E}(G)$ connected together by the very natural property $\pi_e(uv) = \pi_v(u)\pi_v(v)$ for every $uv \in \boldsymbol{E}(G)$. Since we exclusively consider connected graphs with at least one half-edge, the action of an automorphism π on the vertex set is induced by the action of π on half-edges. If G is a simple, then π is determined by the action on the vertices, as is expected. In most situations, we omit subscripts and simply use $\pi(u)$ or $\pi(uv)$. We similarly define *isomorphisms* between different graphs, and we denote existence of an isomorphism by $\cong$. In addition, for colored graphs with three edge types, we require that automorphisms and isomorphisms always preserves the colors, the edge types and the direction of oriented edges.

Automorphism Groups. For undefined concepts and results from permutation group theory, the reader is referred to [26]. We denote groups by Greek letters as for instance Ψ or Γ. We use the following notation for some standard families of groups:

- $\mathbb{S}_n$ for the symmetric group of all n-element permutations,
- $\mathbb{C}_n$ for the cyclic group of integers modulo n,
- $\mathbb{D}_n$ for the dihedral group of the symmetries of a regular n-gon, and
- $\mathbb{A}_n$ for the alternating group of all even n-element permutations.

In this paper, by a group we usually mean a group of automorphisms of a graph acting on the set of half-edges. A group Ψ *acts* on a set S in the following way. Each $g \in \Psi$ permutes the elements of S, and the *action* is described by a mapping $\cdot : \Psi \times S \to S$ where $1 \cdot x = x$ and $(gh) \cdot x = g \cdot (h \cdot x)$. Usually, actions satisfy further properties that arise naturally from the structure of S.

Such an action is the group of all automorphisms of G, denoted by $\mathrm{Aut}(G)$. Each element $\pi \in \mathrm{Aut}(G)$ acts on G, permutes its vertices, edges and half-edges while it preserves edges and incidences between the half-edges and the vertices.

The *orbit* $[v]$ of a vertex $v \in \boldsymbol{V}(G)$ is the set of all vertices $\{\pi(v) \mid \pi \in \Psi\}$, and the orbit $[e]$ of an edge $e \in \boldsymbol{E}(G)$ is defined similarly as $\{\pi(e) \mid \pi \in \Psi\}$. The *stabilizer* of x is the subgroup of all automorphisms which fix x. An action is called *semiregular* if it has no non-trivial (i.e., non-identity) stabilizers of half-edges and vertices. Further, we require the stabilizer of an edge in a semiregular action to be trivial, unless it is a halvable edge, when it may contain an involution transposing the two half-edges. We say that a group is *semiregular* if its action is semiregular. Through the paper, the letter Γ is reserved for semiregular subgroups of $\mathrm{Aut}(G)$. We say that $\pi \in \mathrm{Aut}(G)$ is semiregular if the subgroup $\langle \pi \rangle$ is semiregular. (Note that this is equivalent to the fact that π has all cycles of the same length.)

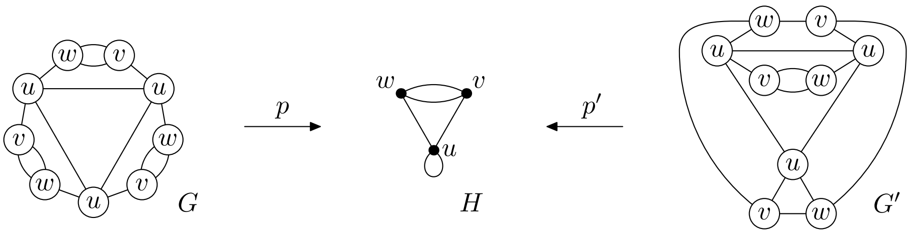

Figure 4: Two covers of H. The projections p_v and p'_v are written inside of the circles, and the projections p_e, p_h, p'_e, and p'_h are omitted. Notice that each loop is realized by having two neighbors labeled the same, and parallel edges are realized by having multiple neighbors labeled the same. Also covering projections preserve degrees.

2.3. Coverings

A graph G *covers* a graph H (or G is a *cover* of H) if there exists a locally bijective homomorphism p called a *covering projection.* A homomorphism p from G to H is given by a mapping $p_h : \boldsymbol{H}(G) \to \boldsymbol{H}(H)$ preserving edges and incidences between half-edges and vertices. It induces two mappings $p_v : \boldsymbol{V}(G) \to \boldsymbol{V}(H)$ and $p_e : \boldsymbol{E}(G) \to \boldsymbol{E}(H)$ such that $p_e(uv) = p_v(u)p_v(v)$ for every $uv \in \boldsymbol{E}(G)$. The property to be local bijective states that for every vertex $u \in \boldsymbol{V}(G)$ the mapping p_h restricted to the half-edges incident with u is a bijection. Figure 4 contains two examples of graph covers. Again, we mostly omit subscripts and just write $p(u)$ or $p(e)$.

Fibers. A *fiber* over a vertex $v \in \boldsymbol{V}(H)$ is the set $p^{-1}(v)$, i.e., the set of all vertices $\boldsymbol{V}(G)$ that are mapped to v, and similarly for fibers over edges and half-edges. We adopt the standard assumption that both G and H are connected. It follows that all fibers of p are of the same size. In other words, $\boldsymbol{h}(G) = k \cdot \boldsymbol{h}(H)$ and $\boldsymbol{v}(G) = k \cdot \boldsymbol{v}(H)$ for some $k \in \mathbb{N}$ which is the size of each fiber, and we say that G is a *k-fold cover* of H.

Regular Coverings. We aim to consider regular coverings which are highly symmetric. From the two examples from Fig. 4, the regular covering p is more symmetric than the non-regular covering p'.

Let Γ be any semiregular subgroup of $\mathrm{Aut}(G)$. It defines a graph G/Γ called a *regular quotient* (or simply *quotient*) of G as follows: The vertices of G/Γ are the orbits of the action Γ on $\boldsymbol{V}(G)$, the half-edges of G/Γ are the orbits of the action Γ on $\boldsymbol{H}(G)$. A vertex-orbit $[v]$ is incident with a half-edge-orbit $[h]$ if and only if the vertices of $[v]$ are incident with the half-edges of $[h]$. (Because the action of Γ is semiregular, each vertex of $[v]$ is incident with exactly one half-edge of $[h]$, so this is well defined.) We say that G *regularly covers* H if there exists a regular quotient of G isomorphic to H.

We naturally construct the *regular covering projection* $p : G \to G/\Gamma$ by mapping the vertices to its vertex-orbits and half-edges to its half-edge-orbits. Concerning an edge $e \in \boldsymbol{E}(G)$, it is mapped to an edge of G/Γ if the two half-edges belong to different half-edge-orbits of Γ. If they belong to the same half-edge-orbits, it corresponds to a half-edge of G/Γ with free end. The projection p is a $|\Gamma|$-fold regular covering.

For the graphs G and H of Fig. 4, we get $H \cong G/\Gamma$ for $\Gamma \cong \mathbb{C}_3$ which "rotates the cycle by three vertices". As a further example, Fig. 5 depicts all quotients of the cube graph.

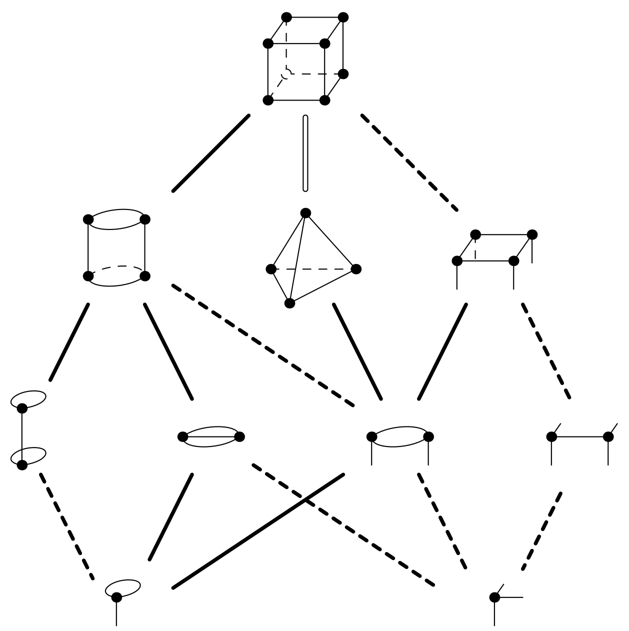

Figure 5: The Hasse diagram of all quotients of the cube graph depicted in a geometric way. When semiregular actions fix edges, the quotients contain half-edges. The quotients connected by bold edges are obtained by 180 degree rotations. The quotients connected by dashed edges are obtained by reflections. The tetrahedron is obtained by the antipodal symmetry of the cube, and its quotient is obtained by a 180 degree rotation with the axis going through the centers of two non-incident edges of the tetrahedron.

2.4. Block-trees and Their Automorphisms

The *block-tree* T of G is a tree defined as follows. Consider all articulations in G and all maximal 2-connected subgraphs which we call *blocks* (with bridge-edges and pendant edges also counted as blocks). The block-tree T is the incidence graph between articulations and blocks. For an example, see Fig. 6. It is well known that every automorphism $\pi \in \mathrm{Aut}(G)$ induces an automorphism $\pi' \in \mathrm{Aut}(T)$.

The Central Block. For a tree, its *center* is either the central vertex or the central pair of vertices of a longest path, depending on the parity of its length. For the block-tree T, all leaves are blocks and each longest path is of an even length. Therefore, T has a central vertex which is either a *central articulation*, or a *central block* of G.

Lemma 2.1. *If G has a non-trivial semiregular automorphism, then G has a central block.*

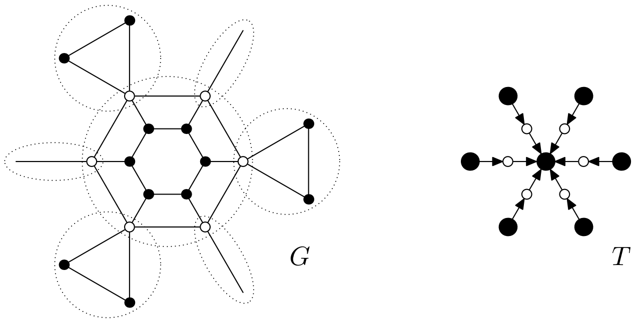

Figure 6: On the left, an example graph G with denoted blocks. On the right, the corresponding block-tree T is depicted, rooted at the central block. The white vertices correspond to the articulations and the big black vertices correspond to the blocks.

Proof. For contradiction, suppose that G has a central articulation u. Every automorphism of a tree preserves its center, so $\mathrm{Aut}(T)$ preserves u. Also, all automorphisms of $\mathrm{Aut}(G)$ preserve u since every automorphism of $\mathrm{Aut}(G)$ induces an automorphism of $\mathrm{Aut}(T)$. This contradicts existence of a non-trivial semiregular automorphism. $\square$

In the following, we shall assume that G contains a central block. We orient the edges of the block-tree T towards the central block; so the block-tree becomes rooted. A *(rooted) subtree* of the block-tree is defined by any vertex different from the central block acting as *root* and by all its predecessors.

Let u be an articulation contained in the central block. By T_u we denote the subtree of T defined by u and all its predecessors, and let G_u be the graph induced by all vertices of the blocks of T_u.

Lemma 2.2. *Let Γ be a semiregular subgroup of $\mathrm{Aut}(G)$. If u and v are two articulations of the central block and of the same orbit of Γ, then $G_u \cong G_v$. Moreover there is a unique $\pi \in \Gamma$ which maps G_u to G_v.*

Proof. Notice that either $G_u = G_v$, or $G_u \cap G_v = \emptyset$. Since u and v are in the same orbit of Γ, there exists $\pi \in \Gamma$ such that $\pi(u) = v$. Consequently $\pi(G_u) = G_v$. Suppose that there exist $\pi, \sigma \in \Gamma$ such that $\pi(G_u) = \sigma(G_u) = G_v$. Then $\pi \cdot \sigma^{-1}$ is an automorphism of Γ fixing u. Since Γ is semiregular, $\pi = \sigma$. $\square$

In the language of quotients, it means that G/Γ consists of C/Γ together with the graphs G_u attached to C/Γ, one for each orbit of Γ.

Why Not Just 2-connected Graphs? Since the behaviour of regular covering with respect to 1-cuts in G is very simple, a natural question follows: why do we not restrict ourselves to 2-connected graphs G? The issue is that the quotient C/Γ might not be 2-connected (see Fig. 12 on the right), so it may consists of many blocks in H. When H contains subtree of blocks isomorphic to G_u, it may correspond to G_u, or it may correspond to a quotient of a subgraph C/Γ, together with some other G_v attached. Therefore, we work with 1-cuts together with 2-cuts and we define 3-connected reduction for 1-cuts in G as well, unlike in [28, 29, 18, 10, 30, 2]. This is essential for the algorithm for regular covering testing described in [13, 14].

3. Structural Properties of Atoms

In this section, we introduce special inclusion-minimal subgraphs of G called atoms. We investigate their structural properties, in particular their behaviour with respect to regular covering projections.

3.1. Definition and Basic Properties of Atoms

Let B be one block of G, so B is a 2-connected graph. Two vertices u and v form a *2-cut* $U = \{u, v\}$ if $B \setminus U$ is disconnected. We say that a 2-cut U is *non-trivial* if $\deg_B(u) \geq 3$ and $\deg_B(v) \geq 3$.

Lemma 3.1. *Let U be a 2-cut and let C be a component of $B \setminus U$. Then U is uniquely determined by C.*

Proof. If C is a component of $B \setminus U$, then U has to be the set of all neighbors of C in B. Otherwise B would not be 2-connected, or C would not be a component of $B \setminus U$. $\square$

The Definition. We first define a set $\mathcal{P}$ of subgraphs of G called *parts* which are candidates for atoms:

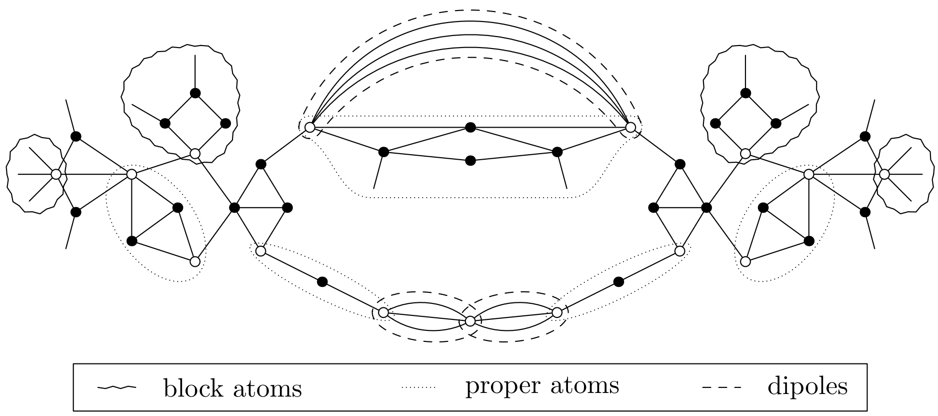

Figure 7: An example of a graph with denoted atoms. The white vertices belong to the boundary of some atom, possibly several of them.

- A *block part* is a subgraph of G induced by the vertices of the blocks of a subtree of the block-tree, non-isomorphic to a pendant edge. (Recall the definitions from Section 2.4.)
- A *proper part* is a subgraph S of G defined by a non-trivial 2-cut U of a block B. The subgraph S consists of a connected component C of $G \setminus U$ (not $B \setminus U$) together with u and v and all edges between $\{u, v\}$ and C. In addition, we require that S does not contain the central block. Therefore, S consists a subgraph of B together with the vertices of the blocks of subtrees of all block-trees attached to C.
- A *dipole part* is any dipole defined as follows. Let u and v be two distinct vertices of degree at least three joined by at least two parallel edges. Then the subgraph induced by u and v is called a *dipole*.

The inclusion-minimal elements of $\mathcal{P}$ are called *atoms*. We distinguish *block atoms*, *proper atoms* and *dipoles* according to the type of the defining part. Block atoms are either stars of pendant edges called *star block atoms*, or pendant blocks possibly with single pendant edges attached to them called *non-star block atoms*. Also each proper atom is a subgraph of a block, together with some single pendant edges attached to it. Notice that a dipole part is by definition always inclusion-minimal, and therefore it is an atom. For an example, see Fig. 7. The above concepts of a proper atom and dipoles have their counter-parts in the literature, they are called pseudo-bricks and bonds, respectively [30]. Some of the following properties and results can be found in literature, see [28, 29, 18, 10, 30] for instance. The novelty of our approach is the use of pendant edges which allow to define atoms also for 1-connected graphs. For readers convenience, we prove all properties to make this paper self-contained.

We use the topological notation to denote the *boundary* ∂A and the *interior* $\mathring{A}$ of an atom A. If A is a dipole, we set $\partial A = \boldsymbol{V}(A)$. If A is a proper or block atom, we put ∂A equal to the set of vertices of A which are incident with an edge not contained in A. For the interior, we use the standard topological definition $\mathring{A} = A \setminus \partial A$ where we only remove the vertices ∂A, the edges adjacent to ∂A are kept in $\mathring{A}$.

Note that $|\partial A| = 1$ for a block atom A, and $|\partial A| = 2$ for a proper atom or dipole A. The interior of a star block atom or a dipole is a set of free edges. Observe for a proper atom A that the vertices of ∂A are exactly the vertices $\{u, v\}$ of the non-trivial 2-cut used in the definition of proper parts. Also the vertices of ∂A of a proper atom are never adjacent in A. Further, no block or proper atom contains parallel edges; otherwise a dipole would be its subgraph, so it would not be inclusion minimal.

Non-overlapping Atoms. Our goal is to replace atoms by edges, and so it is important to know that the atoms cannot overlap too much. The reader can see in Fig. 7 that the atoms only share their boundaries. This is true in general, and we are going to prove it in two steps.

Lemma 3.2. *The interiors of distinct atoms are disjoint.*

Proof. For contradiction, let A and A' be two distinct atoms with non-empty intersections of $\mathring{A}$ and $\mathring{A}'$. First suppose that one of them, say A, is a block atom. Then A corresponds to a subtree of the block-tree which is attached by an articulation u to the rest of the graph. If A' is a block atom then it corresponds to some subtree, and we can derive that $A \subseteq A'$ or $A' \subseteq A$. If A' is a dipole, then it is a subgraph of a block, and thus subgraph of A. If A' is a proper atom, it is defined with respect to some block B. If B belongs to the subtree corresponding to A, then $A' \subseteq A$. Otherwise, a subtree of blocks containing A is attached to A', so $A \subseteq A'$. In both cases, we get contradiction with the minimality. Similarly, if one of the atoms is a dipole, we can easily argue contradiction with the minimality.

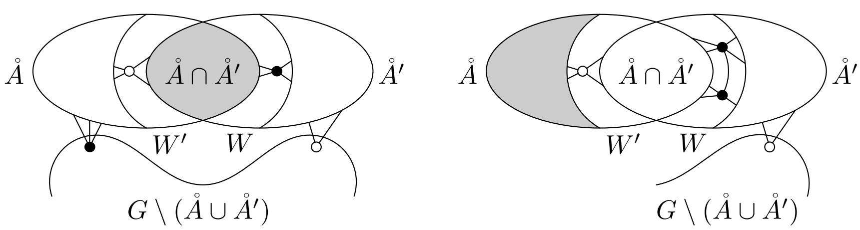

Figure 8: We depict the vertices of ∂A in black and the vertices of $\partial A'$ in white. In both cases, we find a subset of A belonging to $\mathcal{P}$ (its interior is highlighted in gray).

The last case is that both A and A' are proper atoms. Since the interiors are connected and the boundaries are defined as neighbors of the interiors, it follows that both $W' = A \cap \partial A'$ and $W = A' \cap \partial A$ are nonempty. We have two cases according to the sizes of these intersections depicted in Fig. 8.

If $|W| = |W'| = 1$, then $W \cup W'$ is a 2-cut separating $\mathring{A} \cap \mathring{A}'$ which contradicts the minimality of A and A'. If, without loss of generality, $|W| = 2$, then there is no edge between $\mathring{A} \setminus (\mathring{A}' \cup W')$ and the remainder of the graph $G \setminus (\mathring{A} \cup \mathring{A}')$. Therefore, $\mathring{A} \setminus (\mathring{A}' \cup W')$ is separated by a 2-cut W' which again contradicts the minimality of A. We note that in both cases the constructed 2-cut is non-trivial since it is formed by vertices of non-trivial cuts ∂A and $\partial A'$. $\square$

Next we show a stronger version of the previous lemma which states that two atoms can intersect only in their boundaries.

Lemma 3.3. *Let A and A' be two atoms. Then $A \cap A' = \partial A \cap \partial A'$.*

Proof. We already know from Lemma 3.2 that $\mathring{A} \cap \mathring{A}' = \emptyset$. It remains to argue that, say, $\mathring{A} \cap \partial A' = \emptyset$. If A' is a block atom, then $\partial A'$ is the articulation separating A. If A contains this articulation as its interior, it also contains A' as its interior, contradicting $\mathring{A} \cap \mathring{A}' = \emptyset$. Similarly, if A is a block atom, then A' has to be contained in $\mathring{A}$ or vice versa which again contradicts $\mathring{A} \cap \mathring{A}' = \emptyset$.

It remains to argue the case when both A and A' are proper atoms or dipoles. Let $\partial A = \{u, v\}$ and $\partial A' = \{u', v'\}$. First we deal with dipoles. When A is a dipole, it holds since $\mathring{A}$ contains no vertices. If A' is a dipole and A is a proper atom with $u' \in \mathring{A}$, then also the edges of A' belong to A and $A' \subsetneq A$, contradicting the minimality.

We conclude with the remaining case that both A and A' are proper atoms. Recall that ∂A is defined as neighbors of $\mathring{A}$ in G, and that $\partial A'$ are neighbors of $\mathring{A}'$ in G. The proof is illustrated in Fig. 9.

Suppose for contradiction that $\mathring{A} \cap \partial A' \neq \emptyset$ and let $u' \in \mathring{A}$. By definition, u' has at least one neighbor in $\mathring{A}'$, and since $\mathring{A} \cap \mathring{A}' = \emptyset$, this neighbor does not belong to $\mathring{A}$. Therefore, without loss of generality, we have $u \in \mathring{A}'$ and $uu' \in \boldsymbol{E}(G)$. Since A is a proper atom, the set $\{u', v\}$ is not a 2-cut, so there is another neighbor of u in $\mathring{A}$, which has to be equal v'. Symmetrically, u' has another neighbor in $\mathring{A}'$ which is v. So $\partial A \subseteq \mathring{A}'$ and $\partial A' \subseteq \mathring{A}$. If $\partial A = \mathring{A}'$ and $\partial A' = \mathring{A}$, the graph is K_4 (since the minimal degree of cut-vertices is three) which contradicts existence of 2-cuts and atoms. If for example $\mathring{A} \neq \partial A'$, then $\partial A'$ does not cut a subset of $\mathring{A}$, so there exists $w' \in \mathring{A}$ which is a neighbor of $\mathring{A}'$, which contradicts that $\partial A'$ cuts $\mathring{A}'$ from the rest of the graph. $\square$

Primitive Graphs. A graph is called *primitive* if it contains no atoms. The following lemma characterizing primitive graphs can be alternatively obtained from the well-known theorem by Trakhtenbrot [28].[2]

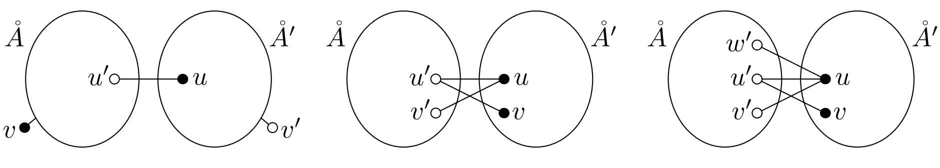

Figure 9: An illustration of the main steps of the proof of Lemma 3.3.

[2]We consider K_1 with an attached single pendant edge as a graph with a central articulation.

Lemma 3.4. *Let G be a primitive graph. If G has a central block, then it is a 3-connected graph, a cycle C_n for $n \geq 2$, or K_2, or can be obtained from the aforementioned graphs by attaching single pendant edges to at least two vertices. If G has a central articulation, then it is K_1, possible with a single pendant edge attached.*

PROOF. The graph G has a central block/articulation. All blocks attached to it have to be single pendant edges, otherwise G would contain a block atom. If G has a central articulation u, after removing all pendant edges, we get a single vertex u, so G is K_1, possibly with a single pendant edge with free ends attached. If G has a central block, after removing all pendant edges, we get the 2-connected graph B consisting of only the central block. We argue that B is one of the stated graphs.

Now, let u be a vertex of the minimum degree in B. If $\deg(u) = 1$, the graph B has to be K_2, otherwise it would not be 2-connected. If $\deg(u) = 2$, then either the graph B is a cycle C_n, or u is an inner vertex of a path connecting two vertices x and y of degree at least three such that all inner vertices are of degree two. But then this path is an atom, a contradiction. Finally, if $\deg(u) \geq 3$, then every 2-cut is non-trivial, and since B contains no atoms, it has to be 3-connected. □

Clearly, the graphs mentioned in the statement are primitive; see Fig. 10.

Structure of Atoms. We call a graph *essentially 3-connected* if it is a 3-connected graph possibly with some single pendant edges attached to it. Similarly, a graph is called *essentially a cycle* if it is a cycle possibly with some single pendant edges attached to it. Similarly to the characterization of primitive graphs in Lemma 3.4, non-star block and proper atoms are either very simple, or almost 3-connected:

Lemma 3.5. *Every non-star block atom A is either K_2 with an attached single pendant edge, essentially a cycle, or essentially 3-connected.*

PROOF. Clearly, the described graphs are possible non-star block atoms. Since A does not contain any smaller block atom, then A is 2-connected graph, possibly with some single pendant edges attached. By removing all single pendant edges, we get a 2-connected graph B, otherwise A contains a smaller block part, which is a smaller block part in G as well.

Let u be a vertex of the minimum degree in B. We have $\deg(u) > 0$, otherwise $B = K_1$ and $A = K_2$. If $\deg(u) = 1$, the graph B has to be K_2, otherwise it would not be 2-connected. If $\deg(u) = 2$, then either the graph B is a cycle C_n, or u is an inner vertex of a path connecting two vertices x and y of degree at least three such that all inner vertices are of degree two. But then this path determines a proper atom in B which is also a proper atom in G, a contradiction. Finally, if $\deg(u) \geq 3$, then every 2-cut is non-trivial, and since B contains no proper atoms, it has to be 3-connected. □

Let A be a proper atom with $\partial A = \{u, v\}$. We define the *extended proper atom* A^+ as A with the additional edge uv.

Lemma 3.6. *For every proper atoms A, the extended proper atom A^+ is either essentially a cycle, or essentially 3-connected.*

PROOF. Clearly, the described graphs are possible extended proper atoms A^+. Notice that A^+ consists a 2-connected graph, possibly with single pendant edges attached, otherwise A contains a smaller block part. By removing all single pendant edges, we get a 2-connected graph B^+, otherwise A^+ contains a smaller block part. Let $\partial A = \{u, v\}$, we have $\deg(u) \geq 2$ and $\deg(v) \geq 2$ in A^+ (and their degrees are preserved in B^+).

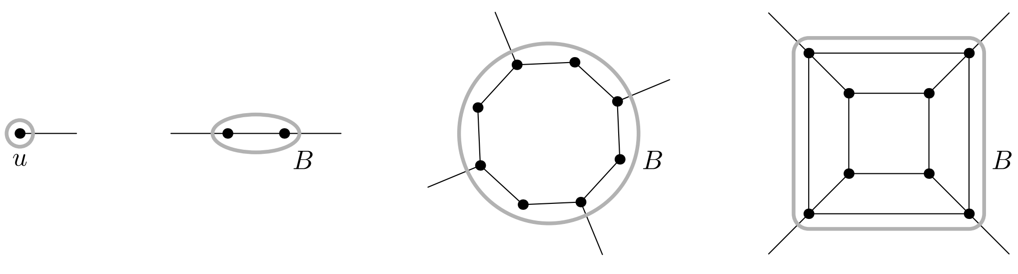

Figure 10: A primitive graph with a central articulation is K_1, and with a central block is either K_2, C_n, or a 3-connected graph, in all these cases with possible single pendant edges attached to it.

Let w be a vertex of the minimum degree in B^+. We have $\deg(w) > 1$, otherwise A again contains a smaller block part. If $\deg(w) = 2$, then either the graph B^+ is a cycle C_n, or w is an inner vertex of a path connecting two vertices x and y of degree at least three such that all inner vertices are of degree two. But then this path is a proper atom in A^+. It corresponds to a proper atom in the original graph since the edge uv in A^+ corresponds to some path in G, so we get a contradiction with the minimality of A. Finally, if $\deg(w) \geq 3$, then every 2-cut is non-trivial, and since B^+ contains no atoms, it has to be 3-connected. □

For all atoms A, single pendant edges are always attached to $\mathring{A}$.

Lemma 3.7. *Let A be an essentially 3-connected graph, and we construct B from A by removing the single pendant edges of A. Then* $\mathrm{Aut}(A)$ *is a subgroup of* $\mathrm{Aut}(B)$.

Proof. These single pendant edges behave like markers, giving a 2-partition of $\boldsymbol{V}(G)$ which $\mathrm{Aut}(A)$ has to preserve. □

3.2. Symmetry Types of Atoms

We distinguish three symmetry types of atoms which describe how symmetric each atom is. For an atom A, we denote by $\mathrm{Aut}(A)$ the setwise stabilizer of ∂A. If A is a block atom, then it is by definition *symmetric*. Let A be a proper atom or dipole with $\partial A = \{u, v\}$. Then we distinguish the following three symmetry types, see Fig. 11:

- *A halvable atom A.* There exits a semiregular involutory automorphism $\tau \in \mathrm{Aut}(A)$ which exchanges u and v.
- *A symmetric atom A.* The atom A is not halvable, but there exists an automorphism in $\mathrm{Aut}(A)$ which exchanges u and v.
- *An asymmetric atom A.* The atom A which is neither halvable, nor symmetric.

When an atom is reduced, we replace it by an edge carrying the type. Therefore we work with multigraphs with three edge types: *halvable edges*, *undirected edges* and *directed edges*. For these multigraphs, we naturally consider only the automorphisms which preserve these edge types and of course the orientation of directed edges, and we use this generalized definition to define symmetry types of their atoms. In the definition of a halvable atom, the automorphism τ fixes no vertices and no directed and undirected edges, but some halvable edges may be fixed.

Action of Automorphisms on Atoms. We show a simple lemma which states how automorphisms behave with respect to atoms.

Lemma 3.8. *Let A be an atom and let $\pi \in \mathrm{Aut}(G)$. Then the following holds:*

(a) The image $\pi(A)$ is an atom isomorphic to A. Further $\pi(\partial A) = \partial\pi(A)$ and $\pi(\mathring{A}) = \mathring{\pi}(A)$.
(b) If $\pi(A) \neq A$, then $\pi(\mathring{A}) \cap \mathring{A} = \emptyset$.
(c) If $\pi(A) \neq A$, then $\pi(A) \cap A = \partial A \cap \partial\pi(A)$.

Proof. (a) Every automorphism permutes the set of articulations and non-trivial 2-cuts. (Recall the definition from the first paragraph of Section 3.1.) So $\pi(\partial A)$ separates $\pi(\mathring{A})$ from the rest of the graph. It follows that $\pi(A)$ is an atom, since otherwise A would not be an atom. And π clearly preserves the boundaries and the interiors.

For the rest, (b) follows from Lemma 3.2 and (c) follows from Lemma 3.3. □

It follows that every automorphism $\pi \in \mathrm{Aut}(G)$ gives a permutation of atoms and $\mathrm{Aut}(G)$ induces an action on the set of all atoms.

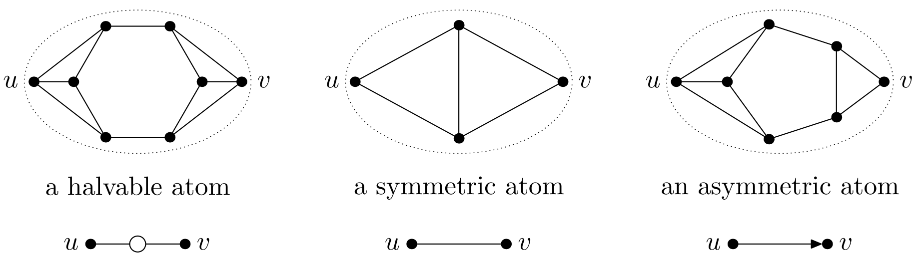

Figure 11: The three types of atoms and the corresponding edge types which we use in the reduction. We denote halvable edges by small circles in the middle.

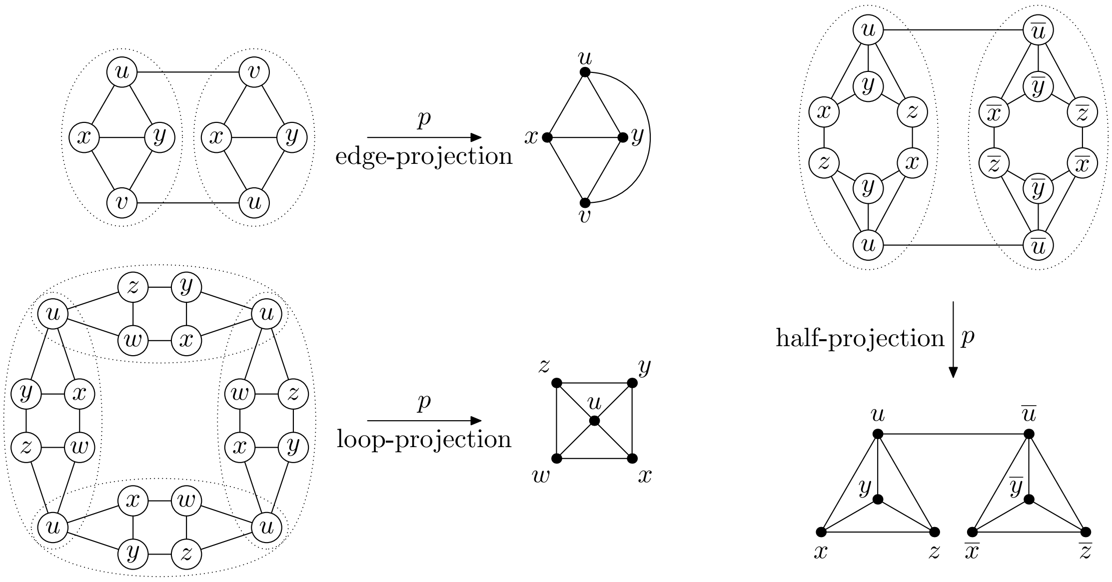

Figure 12: The three cases for projectios of atoms. Notice that for the third graph, an edge-projection can also be applied which gives a different quotient.

3.3. Regular Projections and Quotients of Atoms.

Let Γ be a semiregular subgroup of $\mathrm{Aut}(G)$, which defines a regular covering projection $p : G \to G/\Gamma$. Negami [24, p. 166] investigated possible projections of proper atoms, and we investigate this question in more detail. Let A be an atom and $p|_A$ be its regular projection. We distinguish three different cases illustrated in Fig. 12. For a block atom A, we have exclusively an *edge-projection*. For a proper atom or a dipole A with $\partial A = \{u, v\}$, we get the following three types of projections $p|_A$:

- *An edge-projection.* The atom A is preserved in G/Γ, meaning $p(A) \cong A$. Notice that $p(A)$ is a subgraph of G/Γ, not necessarily induced. For instance for a proper atom, it can happen that $p(u)p(v)$ is adjacent, even through $uv \notin \boldsymbol{E}(G)$, as in Fig. 12.
- *A loop-projection.* The interior $\mathring{A}$ is preserved and the vertices u and v are identified, i.e., $p(\mathring{A}) \cong \mathring{A}$ and $p(u) = p(v)$.
- *A half-projection.* The covering projection p is a $2k$-fold cover. There exists an involutory permutation π in Γ which exchanges u and v and preserves A. Then $p(u) = p(v)$ and $p(A)$ is a halved atom A, consisting of orbits of π on A. This projection can only occur when A is a halvable atom.

Lemma 3.9. *For an atom A and a regular covering projection p, we have $p|_A$ either an edge-projection, a loop-projection, or a half-projection.*

Proof. Let Γ be the semiregular subgroup of $\mathrm{Aut}(G)$ defining $p : G \to G/\Gamma$. For a block atom A, Lemma 2.2 implies that $p(A) \cong A$, so only an edge-projection occurs. It remains to deal with A being a proper atom or a dipole, and let $\partial A = \{u, v\}$. According to Lemma 3.8b every automorphism π either preserves $\mathring{A}$, or $\mathring{A}$ and $\pi(\mathring{A})$ are disjoint.

Suppose that there exists a non-trivial automorphism $\pi \in \Gamma$ preserving $\mathring{A}$. By Lemma 3.8a, we know $\pi(\partial A) = \partial A$, and by semiregularity, π is uniquely determined and exchanges u and v. Then the fiber containing u and v has to be of an even size, with π being an involution reflecting k copies of A, and so p is a $2k$-fold covering projection. Therefore, $p|_A$ is a half-projection.

Suppose that there is no non-trivial automorphism which preserves $\mathring{A}$. The only difference between an edge- and a loop-projection is whether u and v are contained in one fiber of $p|_A$, or not. First, suppose that for every non-trivial $\pi \in \Gamma$ we get $A \cap \pi(A) = \emptyset$. Then no fiber contains more than one vertex of A, and $p_|A$ is an edge-projection, i.e, $A \cong p(A)$. Next, suppose that there exists $\pi \in \Gamma$ such that $A \cap \pi(A) \neq \emptyset$. By Lemma 3.8c, we get $A \cap \pi(A) = \partial A \cap \partial\pi(A)$, so u and v belong to one fiber of $p|_A$, which makes $p|_A$ a loop-projection. $\square$

Figure 13 shows the corresponding quotients $p(A)$ in G/Γ. For an edge-, a loop- and a half-projection $p|_A$, we get three types of quotients $p(A)$ of A which we call an *edge-quotient*, a *loop-quotient* and a *half-quotient*, respectively. The following lemma allows to say “the” edge- and “the” loop-quotient of an atom.

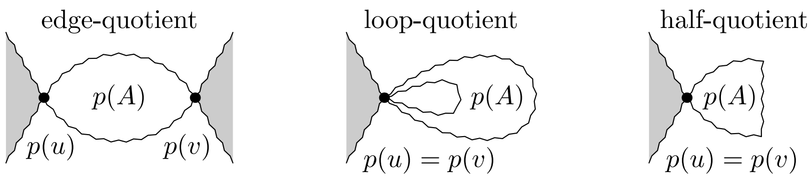

Figure 13: How can the quotient $p(A)$ look in G/Γ, depending on type of $p|_A$.

Lemma 3.10. *For every atom A, there is the unique edge-quotient and the unique loop-quotient up to isomorphism.*

Proof. In both cases, we have $\mathring{A} \cong \mathring{p}(A)$, so these quotients are uniquely determined. $\square$

For half-quotients, this uniqueness does not hold. First, an atom A with $\partial A = \{u, v\}$ has to be halvable to admit a half-quotient. Then each half-quotient is determined by an involutory automorphism τ exchanging u and v; here τ is the restriction of π from the definition of a half-projection. There is a one-to-many relation between non-isomorphic half-quotients and automorphisms τ, i.e., several different automorphisms τ may give the same half-quotient.

Lemma 3.11. *A dipole A has at most $\lfloor \frac{\boldsymbol{e}(A)}{2} \rfloor + 1$ pairwise non-isomorphic half-quotients, and this bound is achieved.*

Proof. Figure 14 shows a construction which achieves the bound. It remains to show that it is an upper bound. Without loss of generality, we can assume that all edges of this dipole are halvable. Let τ be a semiregular involution. Edges which are fixed in τ correspond to half-edges in the half-quotient $A/\langle\tau\rangle$. Pairs of edges interchanged by τ give rise to loops in $A/\langle\tau\rangle$. In the quotient, we have ℓ loops and h half-edges attached to a single vertex such that $2\ell + h = \boldsymbol{e}(A)$. Since ℓ is between 0 and $\lfloor \frac{\boldsymbol{e}(A)}{2} \rfloor$, the upper bound is established. $\square$

For planar proper atoms, we prove in Lemma 5.6 that there are at most two non-isomorphic half-quotients. This non-uniqueness of half-quotients is one of the main algorithmic difficulties for regular covering testing of planar graphs in [13, 14].

4. Graph Reductions and Quotient Expansions

We start with a quick overview. The reduction initiates with a graph G and produces a sequence of graphs $G = G_0, G_1, \ldots, G_r$. To produce G_{i+1} from G_i, we find the collection of all atoms $\mathcal{A}$ in G_i and replace each of them by an edge of the corresponding type. We stop after r steps when a primitive graph G_r containing no further atoms is reached. We call this sequence of graphs starting with G and ending with a primitive graph G_r as the *reduction series* of G.

Suppose that $H_r = G_r/\Gamma_r$ is a quotient of G_r. The reductions applied to reach G_r are reverted on H_r and produce an *expansion series* $H_r, H_{r-1}, \ldots, H_0$ of H_r. We obtain a series of semiregular subgroups $\Gamma_r, \ldots, \Gamma_0$ such that $H_i = G_i/\Gamma_i$ and Γ_i extends Γ_{i+1}. The entire process is depicted in the diagram in Fig. 2.

In this section, we describe structural properties of reductions and expansions. We study changes in automorphism groups done by reductions. Indeed, $\mathrm{Aut}(G_{i+1})$ can differ from $\mathrm{Aut}(G_i)$. But the reduction is done right and important information of $\mathrm{Aut}(G_i)$ is preserved in $\mathrm{Aut}(G_{i+1})$ which is key for expansions. The problem is that expansions are unlike reductions not uniquely determined. From H_{i+1}, we can construct multiple H_i. In this section, we characterize all possible ways how H_i can be constructed from H_{i+1}, and thus establish Theorem 1.2.

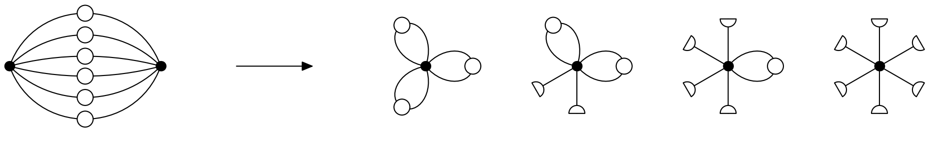

Figure 14: Assuming that quotients can contain half-edges, the depicted dipole has four non-isomorphic half-quotients.

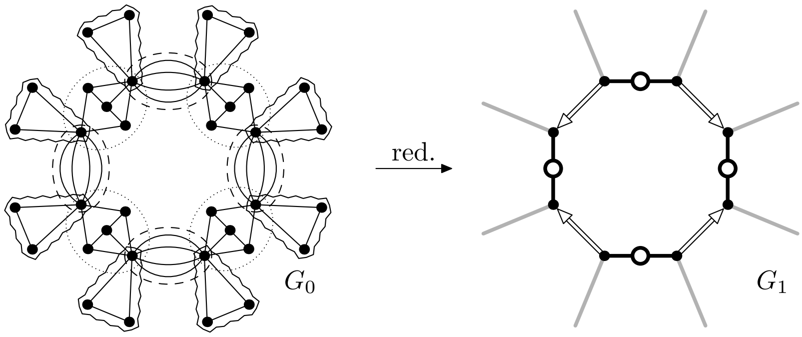

Figure 15: On the left, we have a graph G_0 with three isomorphism classes of atoms, each having four atoms. The dipoles are halvable, the block atoms are symmetric and the proper atoms are asymmetric. We reduce G_0 to G_1 which is an eight cycle with single pendant edges, with four black halvable edges replacing the dipoles, four gray undirected edges replacing the block atoms, and four white directed edges replacing the proper atoms. The reduction series ends with G_1 since it is primitive. Notice the consistent orientation of the directed edges.

4.1. Reducing Graphs Using Atoms

The reduction produces a series of graphs $G = G_0, \ldots, G_r$, by replacing atoms with colored edges encoding isomorphism classes and by edge types encoding symmetry types of atoms.

> **Remark:** *In what follows, we work with multigraphs with colored edges of three types: halvable, undirected and directed. For every automorphism/isomorphism, we require that it preserves colors, edge types and direction of oriented edges.*

We note that the results established in Section 3 transfer to colored graphs and colored atoms without any problems.

For a graph G_i, we find the collection of all atoms $\mathcal{A}$. Two atoms A and A' are isomorphic if there exists an isomorphism which maps ∂A to $\partial A'$. We obtain isomorphism classes for the set of all atoms $\mathcal{A}$ of G_i such that A and A' belong to the same class if and only if $A \cong A'$. To each isomorphism class, we assign one new color not yet used in the graph. The graph G_{i+1} is constructed from G_i by replacing each atom in $\mathcal{A}$ by an edge as follows:

- *A block atom* A with $\partial A = \{u\}$ is replaced by a pendant edge based at u of the color assigned to the isomorphism class containing A.
- *A proper atom or a dipole* A with $\partial A = \{u, v\}$, which is halvable/symmetric/assymmetric, is replaced by a new halvable/undirected/directed edge uv, respectively, of the color assigned to the isomorphism class containing A. It remains to say that for each isomorphism class of asymmetric atom, we consistently choose an arbitrary orientation of the directed edges replacing these atoms.

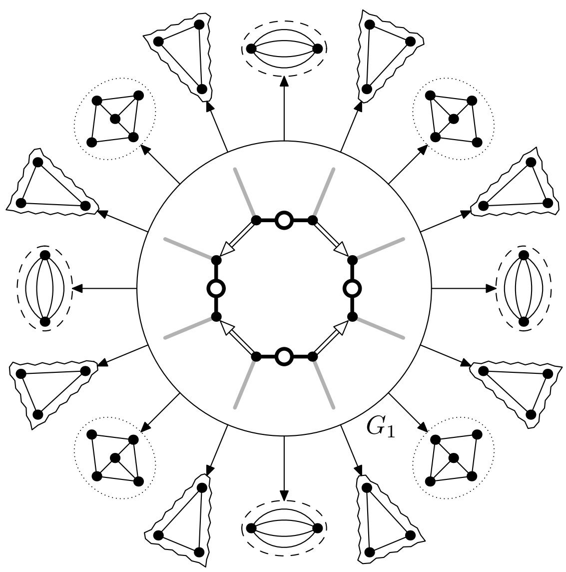

Figure 16: The reduction tree for the reduction series in Fig. 15. The root is the primitive graph G_1 and each leaf corresponds to one atom of G_0.

For an example of the reduction, see Fig. 15.

According to Lemma 3.3, the replaced interiors of the atoms of $\mathcal{A}$ are pairwise disjoint, so the reduction is well defined. We stop in the step r when G_r is a primitive graph containing no atoms. (Recall Lemma 3.4 characterizing all primitive graphs.)

For every graph G, the reduction series corresponds to the *reduction tree* which is a rooted tree defined as follows. The root is the primitive graph G_r, and the other nodes are the atoms obtained during the reductions. If a node contains a colored edge, it has the corresponding atom as a child. Therefore, the leaves are the atoms of G_0, after removing them, the new leaves are the atoms of G_1, and so on. For an example, see Fig. 16.

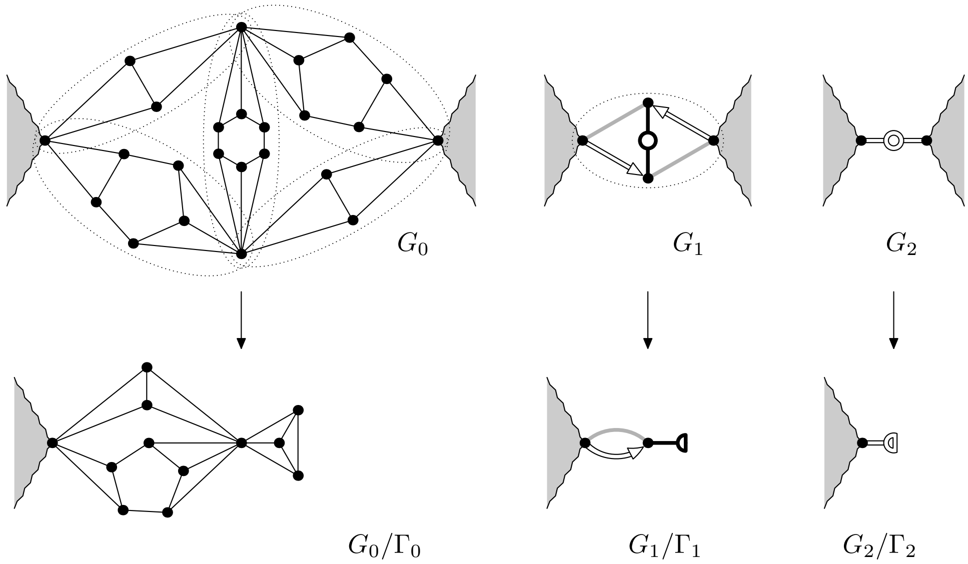

Figure 17: We reduce a part of a graph in two steps. In the first step, we replace five atoms by five edges of different types. As the result we obtain one halvable atom which we further reduce to one halvable edge. Notice that without considering edge types, the resulting atom in G_1 would be just symmetric. In the bottom, we show a part of the corresponding quotient graphs when Γ_i contains a semiregular involutory automorphism π from the definition of a half-projection.

We note that the symmetry type of atoms depends on colors and types of edges the atom contains; see Fig. 17 for an example. Also, the figure depicts a quotient G_2/Γ_2 of G_2, and its expansions to G_1/Γ_1 and G_0/Γ_0. The resulting quotients G_1/Γ_1 and G_2/Γ_2 contain half-edges because Γ_1 and Γ_2 fix some halvable edges but G_0/Γ_0 contains no half-edges. This example shows that in reductions and expansions we need to consider half-edges even both graphs G and G/Γ are simple.

Central Block. We show that the reduction preserves the central block:

Lemma 4.1. *Let G admit a non-trivial semiregular automorphism π. Then each G_{i+1} has a central block which is obtained from the central block of G_i by replacing its atoms by colored edges.*

Proof. By Proposition 4.2c, semiregular automorphisms are preserved during the reduction. By Lemma 2.1, each G_i has a central block. Since we replace only proper atoms and dipoles in the central block, it remains to be a block after the reduction. We argue by induction that it remains central as well.

Let B be the central block of G_i and let B' be this block in G_{i+1}. Consider the subtree T'_u of the block tree T' of G_{i+1} attached to B' in u containing the longest path in T' from B'. This subtree corresponds to T_u in G_i. (See Section 2.4 for the definition of T_u.) Let π be a non-trivial semiregular automorphism in G_i. Then $\pi(u) = v$, and by Lemma 2.2 we have $T_v \cong T_u$. Then T'_v corresponds in G_{i+1} to T_v after reduction and $T'_u \cong T'_v$. Therefore B' is the central block of G_{i+1}. □

If G has a non-trivial semiregular automorphism, then its central block is preserved in the primitive graph G_r. By Lemma 3.4, G_r is either 3-connected, or C_n, or K_2, or can be made from these graphs by attaching single pendant edges to at least two vertices.

We note that in general, the central block/articulation does not have to be preserved and one has to define atoms in all steps of the reduction with respect to the same block/articulation.

Reduction Epimorphism. We describe algebraic properties of the reductions, in particular how the groups $\mathrm{Aut}(G_i)$ and $\mathrm{Aut}(G_{i+1})$ are related. There exists a natural mapping $\Phi_i : \mathrm{Aut}(G_i) \to \mathrm{Aut}(G_{i+1})$ called *reduction epimorphism* which we define as follows. Let $\pi \in \mathrm{Aut}(G_i)$. For the common vertices and edges of G_i and G_{i+1}, we define $\Phi_i(\pi)$ exactly as in π. If A is an atom of G_i, then according to Lemma 3.8a, $\pi(A)$ is an atom isomorphic to A. In G_{i+1}, we replace the interiors of both A and $\pi(A)$ by the edges e_A and $e_{\pi(A)}$ of the same type and color. We define $\Phi_i(\pi)(e_A) = e_{\pi(A)}$. It is easy to see that each $\Phi_i(\pi) \in \mathrm{Aut}(G_{i+1})$.

For purpose of Section 4.2, we also define Φ_i on the half edges. Let $e_A = uv$ and let h_u and h_v be the half-edges composing e_A, and similarly let $h_{\pi(u)}$ and $h_{\pi(v)}$ be the half-edges composing $e_{\pi(A)}$. Then we define $\Phi_i(\pi)(h_u) = h_{\pi(u)}$ and $\Phi_i(\pi)(h_v) = h_{\pi(v)}$.

Proposition 4.2. *The mapping $\Phi_i : \mathrm{Aut}(G_i) \to \mathrm{Aut}(G_{i+1})$ satisfies the following:*

(a) The mapping Φ_i is a group homomorphism.
(b) The mapping Φ_i is an epimorphism, i.e., it is surjective.
(c) For a semiregular subgroup Γ of $\mathrm{Aut}(G_i)$, the restriction $\Phi_i|_\Gamma$ is an isomorphism. Moreover, the subgroup $\Phi_i(\Gamma)$ remains semiregular.

Proof. (a) Clearly, $\Phi_i(\mathrm{id}) = \mathrm{id}$. Let $\pi, \sigma \in \mathrm{Aut}(G_i)$. We need to show that $\Phi_i(\sigma\pi) = \Phi_i(\sigma)\Phi_i(\pi)$. This is clearly true outside the interiors of the atoms. Let A be an atom. By the definition, $\Phi_i(\sigma\pi)$ maps e_A to $e_{\sigma(\pi(A))}$ while $\Phi_i(\pi)$ maps e_A to $e_{\pi(A)}$ and $\Phi(\sigma)$ maps $e_{\pi(A)}$ to $e_{\sigma(\pi(A))}$. So the equality holds everywhere and Φ_i is a group homomorphism.

(b) Let $\pi' \in \mathrm{Aut}(G_{i+1})$, we want to extend π' to $\pi \in \mathrm{Aut}(G_i)$ such that $\Phi_i(\pi) = \pi'$. We just describe this extension on a single edge $e = uv$. If e is an original edge of G, there is nothing to extend. Suppose that e was created in G_{i+1} from an atom A in G_i. Then $\hat{e} = \pi'(e)$ is an edge of the same color and the same type as e, and therefore $\hat{e}$ is constructed from an isomorphic atom $\hat{A}$ of the same symmetry type. The automorphism π' prescribes the action on the boundary ∂A. We need to show that it is possible to define an action on $\mathring{A}$ consistently:

- *A is a block atom:* The edges e and $\hat{e}$ are pendant, attached by articulations u and u'. We define π by an isomorphism σ from A to $\hat{A}$ which takes ∂A to $\partial\hat{A}$.
- *A is an asymmetric proper atom/dipole:* By the definition, the orientation of e and $\hat{e}$ is consistent with respect to π'. Since $\mathring{A}$ is isomorphic to the interior of $\hat{A}$, we define π on $\mathring{A}$ according to one such isomorphism σ.
- *A is a symmetric/halvable proper atom/dipole:* Let σ be an isomorphism of A and $\hat{A}$. Either σ maps ∂A exactly as π', and then we can use σ for defining π. Or we compose σ with an automorphism of A exchanging the two vertices of ∂A. (We know that such an automorphism exists since A is not assymetric.)

So Φ_i is a surjective mapping.

(c) Recall that the kernel $\mathrm{Ker}(\Phi_i)$ is the set of all π such that $\Phi_i(\pi) = \mathrm{id}$ and it is a normal subgroup of $\mathrm{Aut}(G_i)$. It has the following structure: $\pi \in \mathrm{Ker}(\Phi_i)$ if and only if it fixes everything except for the interiors of the atoms. Further, $\pi(\mathring{A}) = \mathring{\pi}(A)$, so π can non-trivially act only inside the interiors of the atoms.

For any subgroup Γ, the restricted mapping $\Phi_i|_\Gamma$ is a group homomorphism with $\mathrm{Ker}(\Phi_i|_\Gamma) = \mathrm{Ker}(\Phi_i) \cap \Gamma$. If Γ is semiregular, then we show that $\mathrm{Ker}(\Phi_i) \cap \Gamma$ is trivial. We know that G_i contains at least one atom A. The boundary ∂A is fixed by $\mathrm{Ker}(\Phi_i)$, so by semiregularity of Γ the intersection with $\mathrm{Ker}(\Phi_i)$ is trivial. Hence $\Phi_i|_\Gamma$ is an isomorphism.

For the semiregularity of $\Phi_i(\Gamma)$, let $\pi' \in \Phi_i(\Gamma)$. Since $\Phi_i|_\Gamma$ is an isomorphism, there exists the unique $\pi \in \Gamma$ such that $\Phi_i(\pi) = \pi'$. If π' fixes a vertex u, then π fixes u as well, so it is the identity, and $\pi' = \Phi_i(\mathrm{id}) = \mathrm{id}$. If π' only fixes an edge $e = uv$, then π' exchanges u and v. If π also fixes e, this edge is halvable and so π' can fix it as well. Otherwise there is an atom A in G_i replaced by e in G_{i+1}. Then $\pi|_A$ is an involutory semiregular automorphism exchanging u and v, so A is halvable. But then e is a halvable edge, and thus π' is allowed to fix it. □

The above statement is an example of a phenomenon known in permutation group theory. Interiors of atoms behave as *blocks of imprimitivity* in the action of $\mathrm{Aut}(G_i)$. It is well-known that the kernel of the action on the imprimitivity blocks is a normal subgroup of $\mathrm{Aut}(G_i)$.

Now, we are ready to prove Proposition 1.1 which states that $\mathrm{Aut}(G_{i+1}) \cong \mathrm{Aut}(G_i)/\mathrm{Ker}(\Phi_i)$:

Proof (Proposition 1.1). By Proposition 4.2b, Φ_i is surjective, so by the well-known Homomorphism Theorem it follows that $\mathrm{Aut}(G_{i+1}) \cong \mathrm{Aut}(G_i)/\mathrm{Ker}(\Phi_i)$. □

Corollary 4.3. *We have* $\mathrm{Aut}(G_r) = \mathrm{Aut}(G_0)/\mathrm{Ker}(\Phi_{r-1} \circ \Phi_{r-2} \circ \cdots \circ \Phi_0)$.

Proof. We have already proved that $\mathrm{Aut}(G_{i+1}) = \mathrm{Aut}(G_i)/\mathrm{Ker}(\Phi_i)$. This equality easily follows from group theory. □

We can also describe the structure of $\mathrm{Ker}(\Phi_i)$:

Lemma 4.4. *The group* $\mathrm{Ker}(\Phi_i)$ *is the direct product* $\prod_{A \in \mathcal{A}} \mathrm{Fix}(A)$ *where* $\mathrm{Fix}(A)$ *is the point-wise stabilizer of* $G_i \setminus \mathring{A}$ *in* $\mathrm{Aut}(G_i)$.

Proof. According to Lemma 3.2, the interiors of the atoms are pairwise disjoint, so $\mathrm{Ker}(\Phi_i)$ acts independently on each interior. Thus we get $\mathrm{Ker}(\Phi_i)$ as the direct product of actions on each interior $\mathring{A}$ which is precisely $\mathrm{Fix}(A)$. $\square$

Alternatively, $\mathrm{Fix}(A)$ can be defined as the point-wise stabilizer of ∂A in $\mathrm{Aut}(A)$. Let $A_1, \dots, A_s$ be pairwise non-isomorphic atoms in G_i, appearing with multiplicities $m_1, \dots, m_s$. According to Lemma 4.4, we get

$$\mathrm{Ker}(\Phi_i) \cong \mathrm{Fix}(A_1)^{m_1} \times \cdots \times \mathrm{Fix}(A_s)^{m_s}.$$

For the example of Fig. 15, we have $\mathrm{Ker}(\Phi_0) \cong \mathbb{C}_2^8 \times \mathbb{C}_2^4 \times \mathbb{S}_4^4$. For the example in Fig. 15, it is shown in [20] that

$$\mathrm{Aut}(G_1) \cong \mathbb{C}_2^2 \qquad \text{and} \qquad \mathrm{Aut}(G_0) \cong (\mathbb{C}_2^8 \times \mathbb{C}_2^4 \times \mathbb{S}_4^4) \rtimes \mathbb{C}_2^2.$$

4.2. Quotients and Their Expansion

Let $G_0, \dots, G_r$ be the reduction series of G and let Γ_0 be a semiregular subgroup of $\mathrm{Aut}(G_0)$. By repeated application of Proposition 4.2c, we get the uniquely determined semiregular subgroups $\Gamma_1, \dots, \Gamma_r$ of $\mathrm{Aut}(G_1), \dots, \mathrm{Aut}(G_r)$ such that $\Gamma_{i+1} = \Phi_i(\Gamma_i)$, each isomorphic to Γ_0. Let $H_i = G_i/\Gamma_i$ be the quotients with preserved colors of edges, and let p_i be the corresponding covering projection from G_i to H_i. Recall that H_i can contain edges, loops and half-edges; depending on the action of Γ_i on the half-edges corresponding to the edges of G_i.

Lemma 4.5. *Every semiregular subgroup Γ_i of $\mathrm{Aut}(G_i)$ corresponds to a unique semiregular subgroup Γ_{i+1} of $\mathrm{Aut}(G_{i+1})$ such that $\Gamma_{i+1} = \Phi_i(\Gamma_i)$.* $\square$

Quotients Reductions. Consider $H_i = G_i/\Gamma_i$ and $H_{i+1} = G_{i+1}/\Gamma_{i+1}$. We investigate relations between these quotients. Let A be an atom of G_i represented by a colored edge e in G_{i+1}. According to Lemma 3.9, $p_i|_A$ is either an edge, a loop or a half-projection. It is easy to see that Φ_i is defined exactly in the way that $p_{i+1}(e)$ corresponds to an edge for an edge-projection, to a loop for a loop-projection, and to a half-edge for a half-projection. (This explains these names of projections and quotients.) Figure 18 shows examples. We get the following commuting diagram:

$$\begin{array}{ccc} G_i & \xrightarrow{\text{red.}} & G_{i+1} \\ \Gamma_i\big\downarrow & & \big\downarrow\Gamma_{i+1} \\ H_i & \xrightarrow{\text{red.}} & H_{i+1} \end{array} \tag{1}$$

So we can construct the graph H_{i+1} from H_i by replacing the projections of atoms in H_i by the corresponding projections of the edges replacing the atoms.

Overview of Quotients Expansions. Our goal is to reverse the horizontal edges in Diagram (1), i.e, to understand:

$$\begin{array}{ccc} G_i & \xleftarrow{\text{exp.}} & G_{i+1} \\ \Gamma_i\big\downarrow & & \big\downarrow\Gamma_{i+1} \\ H_i & \xleftarrow{\text{exp.}} & H_{i+1} \end{array} \tag{2}$$

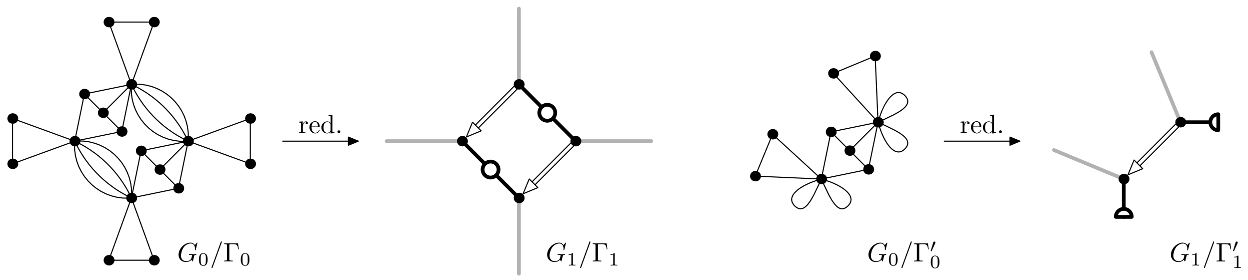

Figure 18: An example of two quotients of the graph G_0 from Fig. 15 with the corresponding quotients of the reduced graph G_1. Here $\Gamma_1 = \Phi_1(\Gamma_0)$ and $\Gamma'_1 = \Phi_1(\Gamma'_0)$.

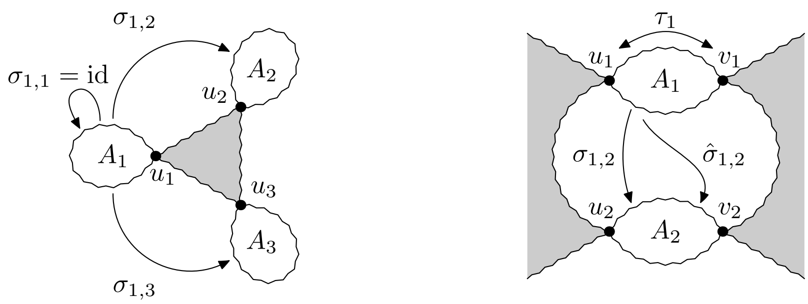

Figure 19: Case 1 is depicted on the left for three edges corresponding to isomorphic block atoms A_1, A_2 and A_3. The depicted isomorphisms are used to extend Γ_{i+1} on the interiors of these atoms. Case 3 is on the right, with an additional semiregular involution τ_1 which transposes u_1 and v_1.

Let Γ_i and Γ_{i+1} be semiregular groups such that $\Phi_i(\Gamma_i) = \Gamma_{i+1}$. Then we call Γ_{i+1} a *reduction* of Γ_i, and Γ_i an *extension* of Γ_{i+1}. There are two fundamental questions we address in this section in full detail:

- *Question 1.* Given a group Γ_{i+1}, which semiregular groups Γ_i are its extensions? Notice that all these groups Γ_i are isomorphic to Γ_{i+1} as abstract groups, but they correspond to different actions on G_i.
- *Question 2.* Let Γ_i and Γ'_i be two semiregular groups extending Γ_{i+1}. Under which conditions are the quotients $H_i = G_i/\Gamma_i$ and $H'_i = G_i/\Gamma'_i$ different?

Extensions of Group Actions. We first deal with Question 1.

Lemma 4.6. *For every semiregular group $\Gamma_{i+1} \le \mathrm{Aut}(G_{i+1})$, there exists an extension $\Gamma_i \le \mathrm{Aut}(G_i)$ such that Diagram (1) commutes.*

Proof. First notice that Γ_{i+1} determines the action of Γ_i everywhere on G_i except for the interiors of the atoms of G_i, so we just need to define it there. Let $e = uv$ be one edge of G_{i+1} replacing an atom A in G_i. Let $|\Gamma_{i+1}| = k$. We distinguish three cases, see Fig. 19:

Case 1: The atom A is a block atom. Then the orbit $[e]$ contains exactly k edges. Let $\partial A = \{u\}$, $[e] = \{e_1, \dots, e_k\}$, and $u_i = \pi'(u)$ for the unique $\pi' \in \Gamma_{i+1}$ mapping e to e_i. (We know that π' is unique because Γ_{i+1} is semiregular.) Let $A_1, \dots, A_k$ be the atoms of G_i corresponding to $e_1, \dots, e_k$ in G_{i+1}. The edges $e_1, \dots, e_k$ have the same color and type, and thus the block atoms A_i are pairwise isomorphic.

We define the action of Γ_i on the interiors of $A_1, \dots, A_k$ as follows. We choose arbitrarily isomorphisms $\sigma_{1,i}$ from A_1 to A_i such that $\sigma_{1,i}(u_1) = u_i$, and put $\sigma_{1,1} = \mathrm{id}$ and $\sigma_{i,j} = \sigma_{1,j}\sigma_{1,i}^{-1}$. If $\pi'(e_i) = e_j$, we set $\pi|_{\mathring{A}_i} = \sigma_{i,j}|_{\mathring{A}_i}$. Since

$$\sigma_{i,k} = \sigma_{j,k}\sigma_{i,j}, \qquad \forall i,j,k, \tag{3}$$

the composition of the extensions π_1 and π_2 of π'_1 and π'_2 is defined on the interiors of $A_1, \dots, A_\ell$ exactly as the extension of $\pi_2\pi_1$. Also, by (3), an identity $\pi'_k\pi'_{k-1}\cdots\pi'_1 = \mathrm{id}$ extends to the identity. Hence the extended action remains semiregular.

Case 2: The atom A is a proper atom or a dipole and the orbit $[e]$ contains exactly k edges. Let $e = uv$ and $[e] = \{e_1, \dots, e_k\}$. We define $u_i = \pi'(u)$ and $v_i = \pi'(v)$ similarly as above. The rest of the argument is similar as in Case 1, we just require that $\sigma_{1,i}(u_1) = u_i$ and $\sigma_{1,i}(v_1) = v_i$.

Case 3: The atom A is a proper atom or a dipole and the orbit $[e]$ contains exactly $\ell = \frac{k}{2}$ edges. Then we have k half-edges in one orbit, so in H_{i+1} we get one half-edge. Let $[e] = \{e_1, \dots, e_\ell\}$. They have to be halvable, and consequently the corresponding atoms $A_1, \dots, A_\ell$ are halvable. Let u_i be an arbitrary endpoint of e_i and let v_i be the second endpoint of e_i. Again, we arbitrarily choose isomorphisms $\sigma_{1,i}$ from A_1 to A_i such that $\sigma_{1,i}(u_1) = u_i$ and $\sigma_{1,i}(v_1) = v_i$, and define $\sigma_{i,j} = \sigma_{1,j}\sigma_{1,i}^{-1}$.

Since A_1 is a halvable atom, we further consider an involution τ_1 of A_1 which exchanges u_1 and v_1. Then τ_1 defines an involution of A_i by conjugation as $\tau_i = \sigma_{1,i}\tau_1\sigma_{1,i}^{-1}$. It follows that

$$\tau_j = \sigma_{i,j}\tau_i\sigma_{i,j}^{-1}, \qquad \text{and consequently} \qquad \sigma_{i,j}\tau_i = \tau_j\sigma_{i,j}, \qquad \forall i,j.$$

We put $\hat{\sigma}_{i,j} = \sigma_{i,j}\tau_i = \tau_j\sigma_{i,j}$ which is an isomorphism mapping A_i to A_j such that $\hat{\sigma}_{i,j}(u_i) = v_j$ and $\hat{\sigma}_{i,j}(v_i) = u_j$. In the extension, we put $\pi|_{\mathring{A}_i} = \sigma_{i,j}|_{\mathring{A}_i}$ if $\pi'(u_i) = u_j$, and $\pi|_{\mathring{A}_i} = \hat{\sigma}_{i,j}|_{\mathring{A}_i}$ if $\pi'(u_i) = v_j$.

Aside (3), we get the following additional identities:

$$\hat{\sigma}_{i,k} = \sigma_{j,k}\hat{\sigma}_{i,j}, \qquad \hat{\sigma}_{i,k} = \hat{\sigma}_{j,k}\sigma_{i,j}, \quad \text{and} \quad \sigma_{i,k} = \hat{\sigma}_{j,k}\hat{\sigma}_{i,j}, \qquad \forall i,j,k. \tag{4}$$

We just argue the last identity:

$$\hat{\sigma}_{j,k}\hat{\sigma}_{i,j} = \tau_k(\sigma_{j,k}\sigma_{i,j})\tau_i = \tau_k\sigma_{i,k}\tau_i = \tau_k\tau_k\sigma_{i,k} = \sigma_{i,k},$$

where the last equality holds since τ_k is an involution. It follows that the composition $\pi_2\pi_1$ is correctly defined as above, and it maps identities to identities.

To conclude the proof, it is easy to observe that by semiregularity of Γ_{i+1} the constructed group Γ_i acts semiregularly on G_i, as well. □

Corollary 4.7. *The construction in the proof of Lemma 4.6 gives all possible extensions of Γ_{i+1} such that Diagram (1) commutes.*

Proof. Recall that the atoms of G_i form blocks in the action of any extension Γ_i, and the action on the blocks is prescribed by Γ_{i+1}. An extension Γ_i of Γ_{i+1} gives in Cases 1 and 2 the isomorphisms $\sigma_{1,i}$, for $i = 1, \dots k$, and in Case 3 the isomorphisms $\sigma_{1,i}$, $\hat{\sigma}_{1,i}$ and τ_i, for $i = 1, \dots \frac{k}{2}$. However, by semiregularity, given these isomorphisms the extension is uniquely determined. □

Quotient Expansion. Recall the description of quotients of atoms from Section 3.3. We are ready to establish the main theorem of this paper. It states that every quotient H_i of G_i can be created from some quotient H_{i+1} of G_{i+1} by replacing edges, loops and half-edges of atoms replaced in the reduction from G_i to G_{i+1} with corresponding edge-, loop- and half-quotients.

Proof (Theorem 1.2). Let $H_{i+1} = G_{i+1}/\Gamma_{i+1}$ and let H_i be constructed in the above way. We first argue that H_i is a quotient of G_i, i.e., it is equal to G_i/Γ_i for some Γ_i extending Γ_{i+1}. To see this, it is enough to construct Γ_i in the way described in the proof of Lemma 4.6. We choose $\sigma_{1,i}$ arbitrarily, and the involutory permutations τ are prescribed by chosen half-quotients replacing half-edges. It is easy to see that the resulting graph is the constructed H_i. We note that only the choices of τ matter, for arbitrary choices of $\sigma_{1,i}$ we get isomorphic quotients H_i.

On the other hand, if H_i is a quotient, it replaces the edges, loops and half-edges of H_{i+1} by some quotients, so we can generate H_i in this way. The reason is that according to Corollary 4.7, we can generate all Γ_i extending Γ_{i+1} by some choices $\sigma_{1,i}$ and τ. □

We say that two quotients H_i and H_i' extending H_{i+1} are *different* if there exists no isomorphism of H_i and H_i' which fixes the vertices and edges common with H_{i+1}. (But H_i and H_i' still might be isomorphic.) According to Lemma 3.10, the edge and loop-quotients are uniquely determined, so we are only free in choosing half-quotients. For non-isomorphic choices of half-quotients, we get different graphs H_i. For instance suppose that H_{i+1} contains a half-edge corresponding to the dipole from Fig. 20. Then in H_i we can replace this half-edge by one of the four possible half-quotients of this dipole.

Corollary 4.8. *If H_{i+1} contains no half-edge, then H_i is uniquely determined. Thus, for an odd order of Γ_r, the quotient H_r uniquely determines H_0.*

Proof. This is implied by Theorem 1.2 and Lemma 3.10 which states that edge- and loop-quotients are uniquely determined. If the order of Γ_r is odd, no half-edges are constructed in H_r, so no half-quotients ever appear. □

Half-quotients of Dipoles. In Lemma 3.11, we describe that a dipole A without colored edges can have at most $\left\lfloor \frac{\boldsymbol{e}(A)}{2} \right\rfloor + 1$ pairwise non-isomorphic half-quotients. This statement can be easily altered to dipoles with colored edges which admit a much larger number of half-quotients:

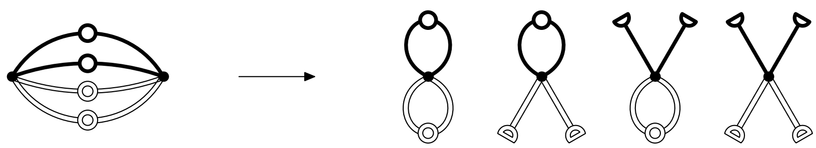

Figure 20: An example of a dipole with a pair of black halvable edges and a pair of white halvable edges. There exist four pairwise non-isomorphic half-quotiens.

Lemma 4.9. *Let A be a dipole with colored edges. Then the number of pairwise non-isomorphic half-quotients is bounded by $2^{\lfloor \frac{\boldsymbol{e}(A)}{2} \rfloor}$ and this bound is achieved.*

Proof. Figure 20 shows an example. It can be easily generalized to exponentially many pairwise non-isomorphic quotients by introducing more pairs of halvable edges of additional colors. It remains to argue correctness of the upper bound.

First, we derive the structure of all involutory semiregular automorphisms τ acting on $\mathring{A}$. We have no freedom concerning the non-halvable edges of A: The undirected edges of each color class has to be paired by τ together. Further, each directed edge has to be paired with a directed edges of the opposite direction and the same color. It remains to describe possible action of τ on the remaining at most $\boldsymbol{e}(A)$ halvable edges of A. These edges belong to c color classes having $m_1, \dots, m_c$ edges. Each automorphism τ has to preserve the color classes, so it acts independently on each class.

We concentrate only on one color class having m_i edges. We bound the number $f(m_i)$ of pairwise non-isomorphic quotients of this class. Then we get the upper bound

$$\prod_{1 \le i \le c} f(m_i) \tag{5}$$

for the number of non-isomorphic half-quotients of A.

The rest of the proof is similar to the proof of Lemma 3.11. An edge e fixed in τ is mapped into a half-edge of the given color in the half-quotient $A/\langle\tau\rangle$. If τ maps e to $e' \neq e$, then we get a loop in the half-quotient $A/\langle\tau\rangle$. The resulting half-quotient only depends on the number of fixed edges and fixed two-cycles in the considered color class. We can construct at most $f(m_i) = \lfloor \frac{m_i}{2} \rfloor + 1$ pairwise non-isomorphic quotients, since we may have zero to $\lfloor \frac{m_i}{2} \rfloor$ loops with the complementing number of half-edges.

The bound (5) is maximized when each class contains exactly two edges. (Except for one class containing either three edges, or one edge if $\boldsymbol{e}(A)$ is odd.) □

Assume that H_{i+1} contains a half-edge corresponding to a half-quotient of a dipole in H_i. By Theorem 1.2, the number of non-isomorphic expansions H_i of H_{i+1} can be exponential in the size difference of H_i and H_{i+1}.

The Block Structure of Quotients. We show how the block structure is preserved during expansions. A block atom A of G_i is always projected by an edge-projection, and so it corresponds to a block atom of H_i. Suppose that A is a proper atom or a dipole, and let $\partial A = \{u, v\}$.

- For an edge-projection $p|_A$, we get $p(u) \neq p(v)$, and $p(A)$ is isomorphic to an atom in H_i.
- For a loop- or a half-projection $p|_A$, we get $p(u) = p(v)$ and $p(u)$ is an articulation of H_i. If A is a dipole, then $p(A)$ is a pendant star of half-edges and loops attached to $p(u)$. By Lemma 3.6, if A is a proper atom, then $p(A)$ is either a path ending with a half-edge and with attached single pendant edges (when A^+ is essentially a cycle), or a pendant block with attached single pendant edges and half-edges (when A^+ is essentially 3-connected). (The reason is that the fiber of an articulation in a 2-fold cover is a 2-cut.)

Lemma 4.10. *The block structure of H_{i+1} is preserved in H_i, possibly with some new subtrees of blocks attached.*

Proof. By Theorem 1.2, edges inside blocks are replaced by edge-quotients of block atoms, proper atoms and dipoles which preserves 2-connectivity. New subtrees of blocks in H_i are created by replacing pendant edges with block atoms, loops by loop-quotients, and half-edges by half-quotients. □

5. Planar Graphs

In this section, we show implication of our theory to planar graphs. Using the reduction, we describe the structure of the automorphism groups of planar graphs. We also characterize the quotients of planar graphs which results in a direct proof of Negami's Theorem. The key point is that regular covering projections behave nicely on 3-connected planar graphs.

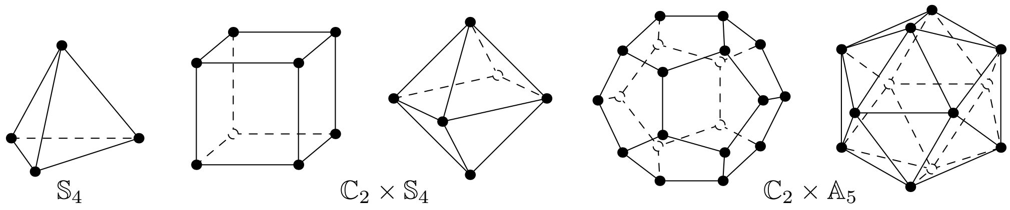

Figure 21: The five platonic solids together with their automorphism groups.

5.1. *Automorphism Groups of 3-connected Planar Graphs*

We review some well-known properties of planar graphs and their automorphism groups. These strong properties are based on Whitney's Theorem [31] stating that a 3-connected graph has a unique embedding into the sphere. This together with the well-known fact that polyhedral graphs are exactly 3-connected planar graphs implies that the automorphism groups of such graphs coincide with the automorphism groups of the associated polyhedra [23].

Spherical Groups. A group is *spherical* if it is the group of the symmetries of a tiling of the sphere. The first class of spherical groups are the subgroups of the automorphism groups of the platonic solids, i.e., $\mathbb{S}_4$ for the tetrahedron, $\mathbb{C}_2 \times \mathbb{S}_4$ for the cube and the octahedron, and $\mathbb{C}_2 \times \mathbb{A}_5$ for the dodecahedron and the icosahedron; see Fig. 21. Table 1 shows the number of conjugacy classes of subgroups of these three groups. Note that conjugate subgroups Γ determine isomorphic quotients G/Γ. The second class of spherical groups is formed by the infinite families $\mathbb{C}_n$, $\mathbb{D}_n$, $\mathbb{C}_n \times \mathbb{C}_2$, and $\mathbb{D}_n \times \mathbb{C}_2$.

Automorphisms of a Map. A map $\mathcal{M}$ is a 2-cell embedding of a graph G onto a surface S. For the purpose of this paper, S is either the sphere or the projective plane. A *rotation* at a vertex is a cyclic ordering of the edges incident with the vertex. When working with abstract maps, they can be viewed as graphs endowed with rotations at every vertex. An *angle* is a triple (v, e, e') where v is a vertex, and e and e' are two incident edges which are consecutive in the rotation at v or in the inverse rotation at v.

An automorphism of a map is an automorphism of the graph which preserves the angles; in other words the rotations are preserved. With the exception of paths and cycles, $\mathrm{Aut}(\mathcal{M})$ is a subgroup of $\mathrm{Aut}(G)$. In general these two groups might be very different. For instance, the star S_n has $\mathrm{Aut}(S_n) = \mathbb{S}_n$, but for any map $\mathcal{M}$ of S_n we just have $\mathrm{Aut}(\mathcal{M}) = \mathbb{D}_n$. If $\mathcal{M}$ is drawn on the sphere, then $\mathrm{Aut}(\mathcal{M})$ is isomorphic to one of the spherical groups [16, 9].

Lemma 5.1 ([31]). *If G is a 3-connected planar graph, then* $\mathrm{Aut}(G)$ *is isomorphic to one of the spherical groups.*

PROOF. Since G is a 3-connected planar graph, there exists the unique embedding of G onto the sphere. Then for any map $\mathcal{M}$ of G, we have $\mathrm{Aut}(G) \cong \mathrm{Aut}(\mathcal{M})$ [31]. □

$\mathbb{S}_4$ of the order 24			
Order	Number	Order	Number
1	1	6	1
2	2	8	1
3	1	12	1
4	3		

$\mathbb{C}_2 \times \mathbb{S}_4$ of the order 48			
Order	Number	Order	Number
1	1	8	7
2	5	12	2
3	1	16	1
4	9	24	3
6	3		

$\mathbb{C}_2 \times \mathbb{A}_5$ of the order 120			
Order	Number	Order	Number
1	1	8	1
2	3	10	3
3	1	12	2
4	3	20	1
5	1	24	1
6	3	60	1

Table 1: The number of conjugacy classes of the subgroups of the groups of platonic solids.

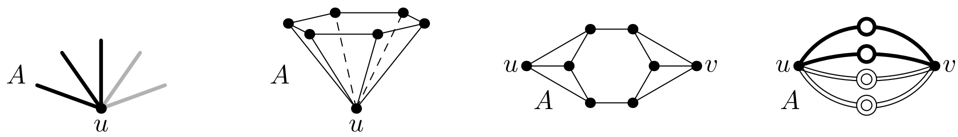

$\mathrm{Fix}(A) \cong \mathbb{S}_2 \times \mathbb{S}_3$	$\mathrm{Fix}(A) \cong \mathbb{D}_6$	$\mathrm{Fix}(A) \cong \mathbb{C}_2$	$\mathrm{Fix}(A) \cong \mathbb{S}_2^2$
$\mathrm{Aut}_{\partial A}(A) \cong \mathbb{S}_2 \times \mathbb{S}_3$	$\mathrm{Aut}_{\partial A}(A) \cong \mathbb{D}_6$	$\mathrm{Aut}_{\partial A}(A) \cong \mathbb{C}_2^2$	$\mathrm{Aut}_{\partial A}(A) \cong \mathbb{S}_2^2 \rtimes \mathbb{C}_2$

Figure 22: An atom A together with its groups $\mathrm{Fix}(A)$ and $\mathrm{Aut}(A)$. From left to right, a star block atom, a non-star block atom, a proper atom, and a dipole.

Primitive Graphs. We start with describing the automorphism groups of primitive graphs.

Lemma 5.2. *For a planar primitive graph G, the group* $\mathrm{Aut}(G)$ *is a spherical group.*

Proof. Recall that a graph is essentially 3-connected if it is a 3-connected graph with attached single pendant edges to some of its vertices. If G is essentially 3-connected, then $\mathrm{Aut}(G)$ is a spherical group from Lemmas 3.7 and 5.1. If G is K_2 or C_n with attached single pendant edges, then it is a subgroup of $\mathbb{C}_2$ or $\mathbb{D}_n$. $\square$

Atoms. Next, we understand the automorphism groups of atoms. For an atom A, recall that $\mathrm{Aut}(A)$ is the set-wise stabilizer of ∂A and $\mathrm{Fix}(A)$ is the point-wise stabilizer of ∂A. See Fig. 22 for examples. The group $\mathrm{Aut}(A)$ is used in Section 5.2 to work with quotients of planar graphs.

Lemma 5.3. *Let A be a planar atom.*

(a) If A is a star block atom, then $\mathrm{Aut}(A) = \mathrm{Fix}(A)$ *which is a direct product of symmetric groups.*
(b) If A is a non-star block atom, then $\mathrm{Aut}(A) = \mathrm{Fix}(A)$ *and it is a subgroup of a dihedral group.*
(c) If A is a proper atom, then $\mathrm{Aut}(A)$ *is a subgroup of* $\mathbb{C}_2^2$ *and* $\mathrm{Fix}(A)$ *is a subgroup of* $\mathbb{C}_2$*.*
(d) If A is a dipole, then $\mathrm{Fix}(A)$ *is a direct product of symmetric groups. If A is symmetric, then* $\mathrm{Aut}(A) = \mathrm{Fix}(A) \rtimes \mathbb{C}_2$*. If A is asymmetric, then* $\mathrm{Aut}(A) = \mathrm{Fix}(A)$*.*

Proof. (a) The edges of each color class of the star block atom A can be arbitrarily permuted, so $\mathrm{Aut}(A) = \mathrm{Fix}(A)$ which is a direct product of symmetric groups.

(b) For the non-star block atom A, $\partial A = \{u\}$ is preserved. We have one vertex in both $\mathrm{Aut}(A)$ and $\mathrm{Fix}(A)$ fixed, thus the groups are the same. Since A is essentially 3-connected, $\mathrm{Aut}(A)$ is a subgroup of $\mathbb{D}_n$ where n is the degree of u.

(c) Let A be a proper atom with $\partial A = \{u, v\}$, and let A^+ be the essentially 3-connected graph created by adding the edge uv. Since $\mathrm{Aut}(A)$ preserves ∂A, we have $\mathrm{Aut}(A) = \mathrm{Aut}(A^+)$, and $\mathrm{Aut}(A^+)$ fixes in addition the edge uv. Because A^+ is essentially 3-connected, $\mathrm{Aut}(A^+)$ corresponds to the stabilizer of uv in $\mathrm{Aut}(\mathcal{M})$ for a map $\mathcal{M}$ of A^+. But such a stabilizer has to be a subgroup of $\mathbb{C}_2^2$. Since $\mathrm{Fix}(A)$ stabilizes the vertices of ∂A, it is a subgroup of $\mathbb{C}_2$.

(d) For an asymmetric dipole, we have $\mathrm{Aut}(A) = \mathrm{Fix}(A)$ which is a direct product of symmetric groups. For a symmetric dipole, we can permute the vertices in ∂A, so we get the semidirect product with $\mathbb{C}_2$. $\square$

5.2. Quotients of Planar Graphs and Negami's Theorem

In this section, we describe quotients of planar graphs geometrically. Using Theorem 1.2, it only remains to understand the quotients of planar primitive graphs and the half-quotients of planar proper atoms. We also show that our structural theory gives a direct proof of Negami's Theorem [24].

Geometry and Quotients. As we have already stated, automorphism groups of 3-connected planar graphs are isomorphic with automorphism groups of the corresponding maps, which allows to use geometry to study regular quotients. We first recall some basic definitions from geometry [27].

Let G be a 3-connected planar graph. An automorphism of G is called *orientation preserving*, if the respective map automorphism preserves the global orientation of the surface. It is called *orientation reversing* if it changes the global orientation of the surface. A subgroup of $\mathrm{Aut}(G)$ is called *orientation preserving* if all its automorphisms are orientation preserving, and *orientation reversing* otherwise. We note that every orientation reversing subgroup contains an orientation preserving subgroup of index two. (The reason is that composition of two orientation reversing automorphisms is an orientation preserving automorphism.)

Let τ be an orientation reversing involution of an orientable surface. The involution τ is called *antipodal* if it is a semiregular automorphism of a closed orientable surface S such that $S/\langle\tau\rangle$ is a non-orientable surface. Otherwise τ is called a *reflection.* A reflection of the sphere fixes a circle. An orientation reversing involution of a 3-connected planar graph is called *antipodal* if the respective map automorphism is antipodal and it is called a *reflection* if the respective map automorphism is a *reflection.* A reflection of a map on the sphere fixes always either an edge, or a vertex.

The quotient of the sphere by an orientation preserving group of automorphisms is again the sphere. The half-quotient of the sphere by a reflection is the disk and and the half-quotient by an antipodal involution is the projective plane. See Fig. 23.

Quotients of Primitive Graphs. By Lemma 3.4, we know that every primitive graph G_r is either 3-connected with attached single pendant edges, or K_2 or C_n with attached single pendant edges. These attached single pendant edges only make $\mathrm{Aut}(G_r)$ smaller, which restricts the possible quotients. Therefore it is sufficient to understand how possible quotients can look for 3-connected planar graphs, K_2 and C_n.

Lemma 5.4 ([27]). *Let G be a 3-connected planar graph and Γ be a semiregular subgroup of $\mathrm{Aut}(G)$. There are three types of quotients of G:*

(a) Rotational quotients – *The action of Γ is orientation preserving and the quotient G/Γ is planar.*
(b) Reflectional quotients – *The action of Γ is orientation reversing but does not contain an antipodal involution. Then the quotient G/Γ is planar and necessarily contains half-edges.*
(c) Antipodal quotients – *The action of Γ is orientation reversing and contains an antipodal involution. Then G/Γ is projective planar.*

Figure 23 shows examples of these types of quotients. We note that an antipodal quotient can be planar, but not necessarily; for an example, see Fig. 1.

The quotients of K_2 are straightforward. Next, we characterize quotients of cycles, which completes the description of possible quotients of primitive graphs:

Lemma 5.5. *Let Γ be a semiregular subgroup of $\mathrm{Aut}(C_n)$. Then C_n/Γ is either a cycle, or a path with two half-edges attached to its ends (only for n even).* □

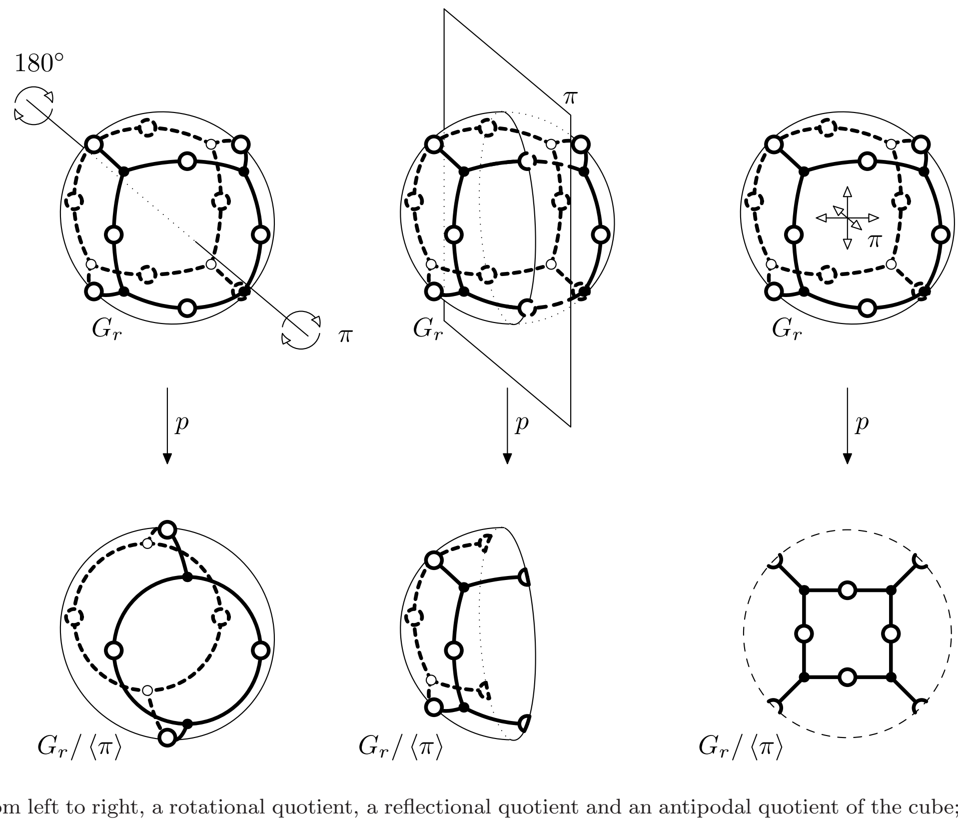

Figure 23: From left to right, a rotational quotient, a reflectional quotient and an antipodal quotient of the cube; also see Fig. 5.

Half-quotients of Proper Atoms. Next, we characterize the half-quotients of planar proper atoms. There are further restrictions compared to the quotients of primitive graphs since the involution has to exchange the vertices of the boundary:

Lemma 5.6. *Let A be a planar proper atom and let $\partial A = \{u, v\}$. There are at most two half-quotients $A/\langle\tau\rangle$ where $\tau \in \mathrm{Aut}(A)$ is an involutory semiregular automorphism transposing u and v:*

(a) The rotational half-quotient – *The involution τ is orientation preserving and $A/\langle\tau\rangle$ is planar with at most one half-edge.*
(b) The reflectional half-quotient – *The involution τ is a reflection and $A/\langle\tau\rangle$ is planar with at least two half-edges.*

Proof. The graph A^+ (obtained from A by adding the edge uv) is an essentially 3-connected planar graph with a unique embedding into the sphere. By Lemma 5.3c, $\mathrm{Aut}(A)$ is a subgroup of $\mathbb{C}_2^2$. An involution τ exchanging u and v corresponds to a map automorphism of A^+ fixing uv. Either τ is a 180° rotation around the centre of uv which gives the rotational half-quotient, or it is a reflection which gives the reflectional half-quotient; see Fig. 24. According to Lemma 5.4, both possible half-quotients are planar. $\square$

Direct Proof of Negami's Theorem. Using the above statements, we give a direct proof of Negami's Theorem. This theorem states that a graph H has a finite planar regular cover G (i.e, $G/\Gamma \cong H$ for some semiregular $\Gamma \leq \mathrm{Aut}(G)$), if and only if H is projective planar. For a given projective planar graph H, the construction of a planar graph G is easy: by embedding H into the projective plane and taking the double cover of this embedding, we get the graph G embedded to the sphere. Below, we prove the harder implication:

Theorem 5.7 (Negami [24]). *Let G be a planar graph. Then every (regular) quotient of G is projective planar.*

Proof. We apply the reduction series on G which produces graphs $G = G_0, G_1, \ldots, G_r$ such that G_r is primitive. If G_r is essentially 3-connected, then by Lemma 5.4 every quotient $H_r = G_r/\Gamma_r$ is projective planar. If G_r is K_2 or C_n with single pendant edges attached, then by Lemma 5.5 every quotient $H_r = G_r/\Gamma_r$ is even planar.

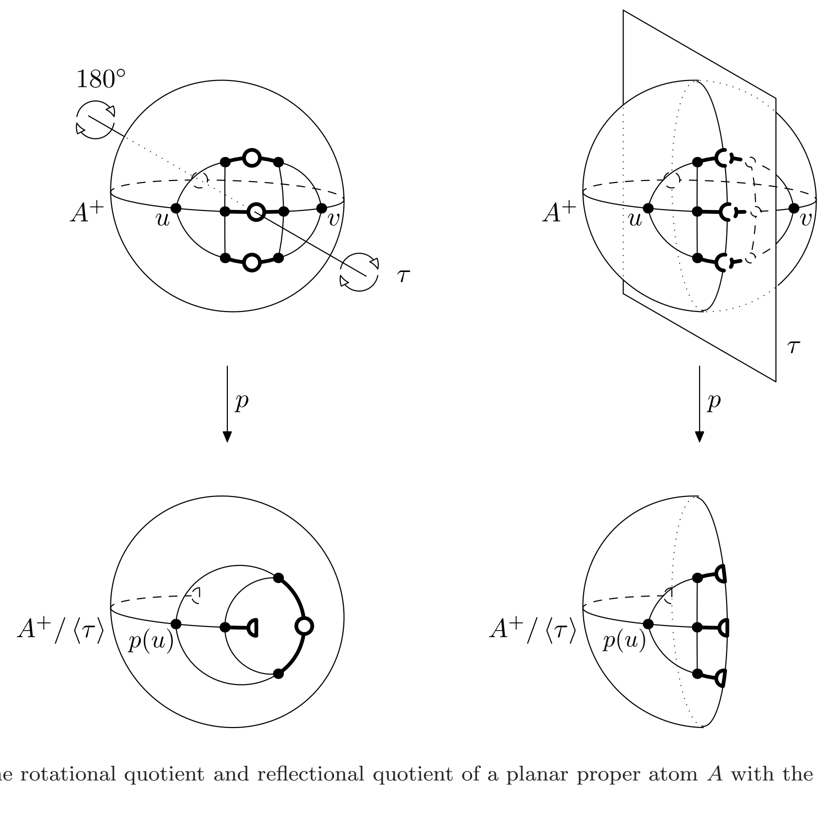

Figure 24: The rotational quotient and reflectional quotient of a planar proper atom A with the added edge uv.

By Theorem 1.2, every quotient $H = G/\Gamma$ can be constructed from some H_r by an expansion series in which we replace edges, loops and half-edges by edge-quotients, loop-quotients and half-quotients, respectively. All edge- and loop-quotients are clearly planar. By Lemma 5.6, every half-quotient of a proper atom is planar, and by Lemma 4.9 every half-quotient of a dipole is a set of loops and half-edges attached to a single vertex, which is also planar. Therefore, these replacements can be done in a way that the underlying surface of H_r is not changed, so H is also projective planar. □

We note that deciding whether $H = G/\Gamma$ is planar or non-planar projective is done on the primitive graph G_r. It is non-planar if and only if Γ_r contains a semiregular antipodal involution and the resulting quotient $H_r = G_r/\Gamma_r$ is non-planar.

6. Concluding Remarks

We recall the main points addressed in this paper:

- We describe the reduction series $G = G_0, \dots, G_r$ such that G_{i+1} is constructed from G_i by replacing the atoms of G_i with colored edges and the primitive graph G_r is either essentially 3-connected, or K_2, or a cycle (Lemma 3.4). We show that $\mathrm{Aut}(G_i)$ is an extension of $\mathrm{Aut}(G_{i+1})$ (Proposition 1.1).
- For a prescribed quotient $H_r = G_r/\Gamma_r$, we describe all possible expansions $H_0 = G_0/\Gamma_0$ which revert the reductions. Theorem 1.2 states that every quotient $H \cong G/\Gamma$ can be obtained in this way, and different quotients H_0 are constructed by non-isomorphic quotients H_r and non-isomorphic choices of half-quotients in the expansions.
- Since the quotients of 3-connected planar graphs can be understood using geometry, we give a direct proof of Negami's Theorem [24] (Theorem 5.7). The reason is that a quotient $H_r = G_r/\Gamma_r$ is due to geometry always planar or projective planar. By Theorem 1.2, the expansions create H from H_r while preserving the underlying surface of H_r.
- Our results have as well algorithmic implications for regular covering testing, described in [13]. In particular, this allows to construct an algorithm for testing whether an input planar graph G regularly covers an input graph H can be constructed, running in time $\mathcal{O}(n^c \cdot 2^{\boldsymbol{e}(H)/2})$.

More General Graphs. Our structural results also work for more general graphs. We have assumed that the graphs G and H are without loops and free half-edges. We can work with loops and half-edges in G in the same way as with pendant edges (of different colors). Since we assume that H contains no half-edges, we set the reductions and expansions in the way that half-edges can appear in the expansion series but no expanded quotient H_0 contains half-edges. This is done by having all edges of G_0 as undirected edges. To admit quotients H_0 with half-edges, it is sufficient to change all edges of G_0 to halvable edges. Also, all the results can be used when G and H contain colored edges, vertices, some edges oriented, etc.

Harmonic Regular Covers. There is a generalization of regular graph covering for which it would be interesting to find out whether our techniques can be modified. Consider geometric regular covers of surfaces, like in Fig. 23 and 24. The orbits of the 180° rotations are of size two, with the exception of two points lying on the axis of the rotation. These exceptional points are called *branch points*. In general, a regular covering projection is locally homeomorphic around a branch point to the complex mapping $z \mapsto z^\ell$ for some integer $\ell \le k$, and ℓ is called the *order* of the branch point.

Assume that G is a 3-connected planar graph embedded onto the sphere, $\Gamma \le \mathrm{Aut}(G)$ is a semiregular subgroup of automorphisms of the sphere, and $p : G \to H = G/\Gamma$ is the regular covering projection. When H is a standard graph (with no free half-edges), all branch points of p belong to faces of the embedding. If a branch point (of order two) is placed in the center of an edge of G, this edge is projected to a half-edge in H. It is possible to consider covering projections between surfaces in which branch points can be placed in vertices of G which gives *harmonic regular covering* [4]. If a branch point of order ℓ is placed in a vertex $v \in \boldsymbol{V}(G)$, then the vertex $p(v) \in \boldsymbol{V}(H)$ has the degree equal $\deg v/\ell$ and for an edge $e \in \boldsymbol{E}(H)$ incident with $p(v)$, the fiber $p^{-1}(e)$ has exactly ℓ edges incident with v.

4-connected Reduction. We have described the way how to reduce a graph to a 3-connected one while preserving its essential structural information. This approach is highly efficient for planar graphs since many problems are much simpler for 3-connected planar graphs; for instance the considered regular graph covering problem. Suppose that we would like to push our results further, say to toroidal or projective planar graphs. The issue is that 3-connectivity does not restrict them much. Is it possible to apply some "4-connected reduction", to reduce the input graphs even further? Suppose that one would generalize proper atoms to be inclusion minimal parts of the graph separated by a 3-cut. Would it be possible to replace them by triangles?

References

[1] J. Abello, M. R. Fellows, and J. C. Stillweil. On the complexity and combinatorics of covering finite complexes. *Australian Journal of Combinatorics*, 4:103–112, 1991.

[2] L. Babai. Automorphism groups of planar graphs II. In *Infinite and finite sets (Proc. Conf. Keszthely, Hungary, 1973) Bolyai-North-Holland*, pages 29–84, 1975.

[3] L. Babai. Automorphism groups, isomorphism, reconstruction. In *Handbook of combinatorics (vol. 2)*, pages 1447–1540. MIT Press, 1996.

[4] M. Baker and S. Norine. Harmonic morphisms and hyperelliptic graphs. *Int. Math. Res. Notes*, 15:2914–2955, 2009.

[5] N. Biggs. *Algebraic graph theory*. Cambridge University Press, 1993.

[6] O. Bílka, J. Jirásek, P. Klavík, M. Tancer, and J. Volec. On the complexity of planar covering of small graphs. In *WG 2011*, volume 6986 of *Lecture Notes in Computer Science*, pages 83–94, 2011.

[7] H. L. Bodlaender. The classification of coverings of processor networks. *Journal of Parallel and Distributed Computing*, 6(1):166–182, 1989.

[8] C. J. Colbourn and K. S. Booth. Linear times automorphism algorithms for trees, interval graphs, and planar graphs. *SIAM J. Comput.*, 10(1):203–225, 1981.

[9] H. S. M. Coxeter. *Introduction to geometry*. New York, London, 1961.

[10] W.H. Cuningham and J. Edmonds. A combinatorial decomposition theory. *Canad. J. Math.*, 32:734–765, 1980.

[11] S. Evdokimov and I. Ponomarenko. Circulant graphs: recognizing and isomorphism testing in polynomial time. *St. Petersburg Mathematical Journal*, 15(6):813–835, 2004.

[12] J. Fiala. Note on the computational complexity of covering regular graphs. In *9th Annual Conference of Doctoral Students, WDS'00*, pages 89–90. Matfyzpress, 2000.

[13] J. Fiala, P. Klavík, J. Kratochvíl, and R. Nedela. Algorithmic aspects of regular graph covers with applications to planar graphs. In *Automata, Languages, and Programming, 41st International Colloquium, ICALP 2014*, volume 8572 of *Lecture Notes in Computer Science*, pages 489–501, 2014.

[14] J. Fiala, P. Klavík, J. Kratochvíl, and R. Nedela. Algorithmic aspects of regular graph covers. *CoRR*, abs/1609.03013, 2016.

[15] J. Fiala and J. Kratochvíl. Locally constrained graph homomorphismsstructure, complexity, and applications. *Computer Science Review*, 2(2):97–111, 2008.

[16] J. L. Gross and T. W. Tucker. *Topological graph theory*. Courier Dover Publications, 2001.

[17] Z. Hedrlín and A. Pultr. On full embeddings of categories of algebras. *Illinois Journal of Mathematics*, 10(3):392–406, 1966.

[18] J. E. Hopcroft and R. E. Tarjan. Dividing a graph into triconnected components. *SIAM Journal on Computing*, 2(3):135–158, 1973.

[19] C. Jordan. Sur les assemblages de lignes. *Journal für die reine und angewandte Mathematik*, 70:185–190, 1869.

[20] P. Klavík, R. Nedela, and P. Zeman. Constructive approach to automorphism groups of planar graphs. abs/1506.06488, 2015.

[21] J. Kratochvíl, A. Proskurowski, and J. A. Telle. Covering regular graphs. *J. Comb. Theory Ser. B*, 71(1):1–16, 1997.

[22] A. Malnič, R. Nedela, and M. Škoviera. Lifting graph automorphisms by voltage assignments. *European Journal of Combinatorics*, 21(7):927–947, 2000.

[23] P. Mani. Automorphismen von polyedrischen graphen. *Math. Ann.*, 192:297–303, 1971.

[24] S. Negami. The spherical genus and virtually planar graphs. *Discrete Mathematics*, 70(2):159–168, 1988.

[25] G. Ringel and J. Youngs. Solution of the Heawood map-coloring problem. *Proceedings of the National Academy of Sciences of the United States of America*, 60(2):438–445, 1968.

[26] J.J. Rotman. *An Introduction to the Theory of Groups*. Graduate Texts in Mathematics. Springer, 1994.

[27] J. Stillwell. *Geometry of surfaces*. Springer, 1992.

[28] B.A. Trakhtenbrot. Towards a theory of non-repeating contact schemes. *Trudi Mat. Inst. Akad. Nauk SSSR*, 51:226–269, 1958.

[29] W. T. Tutte. *Connectivity in graphs*, volume 15. University of Toronto Press, 1966.

[30] T.R.S. Walsh. Counting unlabeled three-connected and homeomorphically irreducible two-connected graphs. *J. Combin. Theory B*, 32:12–32, 1982.

[31] H. Whitney. Nonseparable and planar graphs. *Trans. Amer. Math. Soc.*, 34:339–362, 1932.